


THE COMPLETE COLLECTION

THE HARDWARE NEVER HANDBOOK

Every part, technology, and decision that passes on the bench and fails in the real world, organized by stage, with the fix for each one.



BENCH PROTOTYPE
WORKS ON YOUR DESK

FIX →



PRODUCTION DESIGN
SURVIVES YOUR CUSTOMERS

PREPARED BY
JOHN TEEL

DOC NO.
NVR-01

REV
A

SOURCE
[PREDICTABLEDESIGNS.COM](https://predictabledesigns.com)

Contents

Thirteen parts, ordered by the stage of development where each decision gets made.

Every row in this table is a link, so click a part to jump straight to it. Each part opens with its own list of entries, and every entry there is a link as well. The title in the header of any page brings you back here, and the bookmarks panel in your PDF reader lists every entry.

#	PART	STAGE OF DEVELOPMENT	PAGE
	How to use this handbook	Start here	3
1	Part 1: Microcontrollers and Processors	Architecture and platform selection	4
2	Part 2: Arduino and Development Boards	Prototyping	14
3	Part 3: Wireless Modules and Connectivity	Component selection	21
4	Part 4: Sensors	Component selection	33
5	Part 5: Batteries and Charging	Power architecture	41
6	Part 6: Voltage Regulators and Power Conversion	Power design	52
7	Part 7: Capacitors, Passives, and Discretes	Schematic design	63
8	Part 8: Connectors, Interfaces, and Peripherals	Schematic and industrial design	73
9	Part 9: Circuit Protection	Schematic design	81
10	Part 10: PCB Layout	PCB layout	88
11	Part 11: Firmware	Firmware	100
12	Part 12: Enclosure Design	Mechanical design	110
13	Part 13: Sourcing and Counterfeit Parts	Sourcing and production	120
	Appendix: The Master Swap Table	Quick reference	132
	The Production Filters	Before you sign off	139

The parts run in the order the decisions come up in a real product: the processor and platform first, then the parts around it, then the schematic-level protection, the layout, the firmware, the enclosure, and finally the sourcing decisions that have to hold up at production volume.

How to Use This Handbook

Read the part that matches the decision in front of you, then run the checks before you commit.

Every part, board, technology, and design decision in this handbook works on a bench. That is exactly what makes them dangerous. A prototype only has to survive you, and you always use the right adapter, the right battery, and the right firmware build. A product has to survive thousands of strangers, a test lab, a freight forwarder, a customs inspector, and a second production run a year from now when the cheap part you designed around is no longer available.

This handbook collects every entry from my "You Should Never" video series, plus my video on the technologies to avoid, into one reference. Instead of running video by video, the entries are grouped by the kind of decision they belong to and ordered by the stage of development where you make it. If you are choosing a microcontroller, read Part 1. If your boards are about to go to the fab, read Part 10. If you are sourcing parts for the first production run, read Part 13 before you send anyone money.

Each entry goes deeper than the video it came from. The tables list more alternatives, the advice is more specific, and every entry ends with the checks I would run before committing to a decision. Nothing in here assumes you have a compliance engineer or an RF lab on staff, because most of the product creators I work with do not.

How Each Entry Is Laid Out

SECTION	WHAT IT TELLS YOU
The band at the top	The entry number, what you should never do, and the category it belongs to. The number is part.entry, so entry 6.3 is the third entry in Part 6, and that is how the cross-references in the text work.
Why it fails in a product	The mechanism behind the failure, not just the symptom. Most of these problems never show up on a bench, so this section explains what changes once the product is in a customer's hands.
Swap it for / Do this instead	The specific replacement part, technology, or design change, and the reason it holds up in production. Where the video named one alternative, the table usually names several.
Before you commit	The datasheet specs to verify, the questions to ask a supplier, or the tests to run before this decision is final. Treat these as tick boxes for your own design.
The callout	A cost check, a warning, or a story from the field that puts the entry in perspective.

One rule that runs through the whole handbook

The further a design travels from your bench, the more of this applies. A one-off sitting on your desk can survive a bad part, because you are right there to replace it. A project that has to run unattended for a year cannot, and a product cannot at all, because every decision gets multiplied by every unit you build and it has to be the same decision when you reorder next year. Read each entry against how far yours has to go.

STAGE: ARCHITECTURE AND PLATFORM

MICROCONTROLLERS AND PROCESSORS

Choosing the wrong processor can lock you into months of rework, destroy your margins, or leave you stranded when the manufacturer moves on. After reviewing hundreds of designs, these are the chips I keep seeing in new commercial products where they have no business being, plus the two processor decisions that sink small teams before the first board is even laid out.

IN THIS PART

1.1	ESP8266 in a New Design	p. 5
1.2	PIC16 for a New Product	p. 6
1.3	Ultra-Cheap No-Name Microcontrollers	p. 7
1.4	STM8	p. 8
1.5	MSP430 for a New General-Purpose Product	p. 9
1.6	ATmega328P (the Classic Arduino Chip)	p. 10
1.7	A High-End MCU Where an MPU Is Cheaper	p. 11
1.8	High-End Application Processors	p. 12
1.9	Custom IC Design	p. 13

1.1

ESP8266 in a New Design

WI-FI MICROCONTROLLER

WHY IT FAILS IN A PRODUCT

- It is a single-core chip with limited GPIO, only a basic 10-bit ADC, and an aging SDK that is not seeing meaningful new development. All of Espressif's real momentum is on the ESP32 family now.
- The ESP32-C3 costs nearly the same, so there is no cost argument left for starting a brand-new product on the 8266.
- It has none of the security features a connected product now needs. There is no secure boot and no flash encryption, so the firmware protection covered in Part 11 is not even available to you.
- The real danger is that outdated tutorials and online courses still teach it, so beginners default to it without realizing they are starting on a dead-end platform, and every line of firmware they write has to be ported later.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
Espressif ESP32-C3	This is the closest match in price and simplicity, with a single core and a small footprint, and it adds Bluetooth LE, a much better peripheral set, secure boot, and flash encryption.
Espressif ESP32-C6	If you want Wi-Fi 6, Thread, or Zigbee alongside Bluetooth LE, the C6 adds them for a small premium over the C3.
Espressif ESP32-S3	If the product needs more processing, more GPIO, native USB, or a display, the dual-core S3 is the step up inside the same ecosystem.
A pre-certified WROOM or MINI module	When you move to your own PCB, use a shielded module carrying FCC modular approval and CE test documentation rather than the bare chip, unless you have RF expertise on the team.

BEFORE YOU COMMIT

- **Peripheral count:** List every GPIO, ADC channel, and interface the product needs, add twenty percent for the features you have not thought of yet, and check that count against the C3 pinout.
- **Security plan:** Decide now how secure boot and flash encryption get enabled in production programming, because both are painful to add after the first units ship.
- **Toolchain:** Start on ESP-IDF, or plan the move before the firmware grows, since the production features live there. The Arduino core can run inside an ESP-IDF project, so your code is not thrown away.

If you already ship on the 8266

Espressif's longevity commitment for the 8266 runs until 2029, so an existing product is not an emergency. It is simply a bad starting point for anything new, and the same money buys a far better chip.

1.2 PIC16 for a New Product

8-BIT MICROCONTROLLER

WHY IT FAILS IN A PRODUCT

- It is an 8-bit architecture that has been outclassed for most new product designs, with limited RAM and flash, and a cost advantage over 32-bit parts that has essentially disappeared.
- A PIC16 and an ARM Cortex-M0+ part often cost about the same now, but the Cortex-M0+ gives you dramatically more processing power, memory, and peripheral options for that money.
- Microchip offers mature tools, but the broader ARM ecosystem gives you far more third-party libraries, more community support, and a much bigger pool of developers who already know the platform.
- Product creators usually inherit it through a legacy codebase, or because a contract engineer defaults to the platform they learned a decade ago. That means the decision was made for the engineer's comfort and not for the product.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
ST STM32C0	This is ST's entry-level Cortex-M0+ family priced against 8-bit parts, and it scales to the G0 and G4 families with minimal rework when the product grows.
TI MSPM0	TI's Cortex-M0+ line is aimed at the same sub-dollar bracket, with TI's documentation and application support behind it.
Microchip SAM D or SAM L	If you want to stay inside Microchip's tools and support, their own Cortex-M0+ families give you a modern core without leaving the vendor.

BEFORE YOU COMMIT

- **Price at volume:** Get real quotes for the PIC16 and for an STM32C0 or MSPM0 at your target quantity and put them side by side, because the gap you remember is usually gone.
- **Room to grow:** Estimate the RAM and flash you will need two firmware revisions from now, not just for the first release, and pick a family with a pin-compatible step up.
- **Ask the engineer why:** If a contractor specifies a PIC16, ask what the product gains from it. Familiarity is a fine answer for a prototype and a bad one for a product.

The exception

A product already shipping on a PIC16 with stable firmware is not a reason to redesign. This entry is about new designs, where you are choosing the platform you will live with for years.

1.3

Ultra-Cheap No-Name Microcontrollers

SUB-\$0.10 MICROCONTROLLERS

WHY IT FAILS IN A PRODUCT

- The pitch is a microcontroller for under ten cents, and the reality is machine-translated datasheets, limited or nonexistent English SDK support, a tiny community, and almost no reference designs.
- When something goes wrong, and it almost always does, you need documentation and community support you can actually rely on. I have seen product creators stuck for weeks on a peripheral bug that nobody online has ever encountered, simply because the chip does not have enough users to generate a knowledge base.
- An undocumented silicon bug is the worst kind, because there is no errata sheet telling you it exists, so you spend that time doubting your own firmware.
- Sourcing is unreliable too, since these parts appear and disappear from the market without warning, and a chip you cannot buy from a major distributor is a chip you cannot count on reordering.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
ST STM32C0 or TI MSPM0	Both are cheap, come with mature documentation and large communities, and are stocked by every major distributor, which is what you actually need when a peripheral misbehaves.
WCH CH32V003	This sub-ten-cent RISC-V chip is the rare exception with decent English documentation and a growing community, but exceptions like this are still rare, so treat it as one.
A slightly larger part from the same major vendor	If the reason for the no-name chip was a few cents, a volume quote from ST, TI, or Microchip usually closes most of that gap.

BEFORE YOU COMMIT

- **Community test:** Search for an active forum or community around the exact chip. If nobody is asking questions about it, nobody is answering them either.
- **Documentation test:** Read the getting-started guide and one peripheral chapter of the datasheet. If either one was obviously run through a translator, expect the errata to be missing as well.
- **Reference design test:** Ask for a reference design from the manufacturer, with schematics. If there is none, you are the reference design.

When it can work

These chips can work in an extremely simple product where the firmware is basic and the odds of hitting an undocumented silicon bug are low. Even then, weigh a twenty-cent saving per unit against one week of an engineer's time, because that week costs more than the savings on the first several thousand units.

1.4

STM8

PROPRIETARY 8-BIT

WHY IT FAILS IN A PRODUCT

- ST has made it clear that their future is the STM32, and while the compiler for the STM8 is now free, the architecture itself is proprietary and not ARM-based, so your toolchain skills and your hardware-level code do not transfer to anything else.
- The cost argument used to hold up, but STM32C0 and STM32G0 parts have dropped to price points that overlap with the STM8, and they give you an ARM Cortex-M0+ with a modern ecosystem.
- That ecosystem difference matters more than most people realize, because it means access to more third-party libraries, more example code, and a much larger pool of engineers who can work on your project.
- If you outgrow the STM8, you are rewriting every driver and rebuilding your toolchain. If you start on an STM32C0, you can scale up to the G0, G4, or beyond with far less rework.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
ST STM32C0	It sits at STM8 pricing with a Cortex-M0+ core, and it is the bottom rung of a ladder you can climb without changing toolchains.
ST STM32G0	When you need more peripherals, more memory, or better analog, the G0 is the next step and shares the same tools and libraries.
ST STM32G4	For products that later need real signal processing or motor control, the G4 is a further step on the same path.

BEFORE YOU COMMIT

- **Price at volume:** Compare the STM8 part against the STM32C0 at your quantity. If the difference is a few cents, the ecosystem is worth far more than that.
- **Engineer pool:** Ask who could take over the firmware if your current engineer left. With STM8 that list is short, and with STM32 it is long.
- **Reuse plan:** If existing STM8 code is the reason for staying, estimate the port realistically. Peripheral drivers rarely survive the move, and application logic usually does.

A free compiler is not the point

The STM8 became more attractive when its compiler went free, but the cost that matters is the engineering time locked into an architecture with nowhere to go.

1.5

MSP430 for a New General-Purpose Product ULTRA-LOW-POWER MICROCONTROLLER

WHY IT FAILS IN A PRODUCT

- This one hits close to home since I spent years at TI and I love TI parts, but TI has introduced the MSPM0 family and published detailed migration guides encouraging MSP430 engineers to move to it.
- TI still supports the 430 and has released some newer variants, but for most new general-purpose designs TI's own momentum is moving toward the M0.
- The 430's killer feature was always ultra-low-power performance, but for most new designs the MSPM0 and STM32U0 families have closed much of that gap, and they give you a modern ARM ecosystem on top of it.
- The 430 is a proprietary 16-bit architecture, so the same portability problem as the STM8 applies. The hardware-level code and the toolchain skills stay inside one vendor's legacy family.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
TI MSPM0	This is TI's direct migration path, with the low-power focus of the 430, a Cortex-M0+ core, and TI's own migration guides to move existing code.
ST STM32U0	If you want to stay in ultra-low-power territory but are not tied to TI, the U0 is built for the same battery-first products.
MSP430FR (FRAM) parts	FRAM gives you non-volatile storage with virtually unlimited write endurance and very low power during writes, so for energy harvesting or constant data logging, staying on an FRAM 430 is a legitimate choice.

BEFORE YOU COMMIT

- **Sleep current, measured:** Write down the sleep current and the wake-up time your product actually needs, then compare the MSPM0 and STM32U0 datasheets against those numbers rather than against the 430's reputation.
- **Write endurance:** Count how often the product writes to non-volatile memory over its life. If the answer runs into the millions, FRAM is the reason to stay, and Part 11 explains why.
- **Lifecycle status:** Check TI's product page for the exact 430 part number, since the family is broad and the status varies part by part.

The one place the 430 still wins

FRAM is the exception that proves the rule. If your product logs data continuously or runs on harvested energy, an FRAM-based 430 solves a problem the M0 families do not, and that is a real product reason rather than nostalgia.

1.6

ATmega328P (the Classic Arduino Chip)

8-BIT AVR

WHY IT FAILS IN A PRODUCT

- The Arduino platform is the gateway into microcontrollers for a huge number of people, and it is fantastic for learning and prototyping. The problems start when a founder tries to ship a product on it.
- The ATmega328P is underpowered for the price, has no wireless, and has limited peripherals. At its price point you can get a 32-bit ARM chip with several times more memory, better peripherals, and a much more accessible debug interface.
- The classic Arduino workflow abstracts away everything at the hardware level, which is great for learning and terrible for production. When you need to optimize power, debug a timing issue, or write a custom bootloader, that abstraction fights you instead of helping you.
- Microchip has flagged several ATmega328P variants as not recommended for new designs, and their own development focus has clearly shifted to newer AVR families.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
Espressif ESP32-C3	If the product needs wireless, the C3 gives you Wi-Fi and Bluetooth LE with a modern core for about the same money.
ST STM32C0	If the product does not need wireless, the C0 is cheaper or similarly priced with far more capability and a proper SWD debug interface.
Raspberry Pi RP2040	A dual-core part under a dollar in production quantities, with plenty of I/O and a large community, and you can prototype on a Pico for a few dollars.
Newer AVR families	If the firmware is deeply tied to AVR, Microchip's current AVR DA and DB families are where their development focus went.

BEFORE YOU COMMIT

- **Lifecycle check:** Look up the exact 328P variant on Microchip's site and read its status. Not recommended for new designs means exactly what it says.
- **Real cost:** Price the 328P together with the crystal, regulator, and USB bridge the Arduino board hid from you, then price a C3 or C0 with the same support parts.
- **Debug path:** Confirm you can attach a real debugger with breakpoints to whatever you pick. Serial print debugging does not scale to a product.

Yes, there are products shipping on it

Plenty of products have shipped with ATmega chips inside, but for someone starting a new product today there is almost no scenario where it is the best choice. The Arduino IDE works with every alternative above.

1.7 A High-End MCU Where an MPU Is Cheaper

PROCESSOR ARCHITECTURE

WHY IT FAILS IN A PRODUCT

- Product creators reach for a top-tier STM32H7 or an NXP crossover MCU because they need a display, networking, a complex user interface, or heavy data processing, and staying in the MCU world feels safer, with no operating system, no Linux boot time, and a bare-metal feel.
- But when you are paying 8 to 15 dollars for an MCU and then adding external RAM, external flash, and a display controller, you have likely exceeded the silicon cost of a microprocessor that handles all of that natively.
- You also give up thousands of existing Linux drivers and software packages, which can dramatically cut development time for networking, file systems, and graphics.
- This is not about simple battery-powered products or basic embedded control. It is about products that are already creeping into application-processor territory whether you admit it or not.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Run the rule of thumb	If the design needs external RAM, a display larger than a small OLED, or an RTOS complex enough to basically be an OS, evaluate a microprocessor.
Evaluate the ST STM32MP1	It gives you Linux, display support, and networking at a chip cost that competes with a loaded high-end MCU setup, inside ST's tools.
Evaluate the Allwinner T113 or a Raspberry Pi Compute Module	Both give you Linux and a massive software ecosystem, and the Compute Module adds real documentation and reference designs.
Keep the MCU for hard real time	If the product needs microsecond-level timing for motor control or precision actuation, a high-end MCU is still the right tool, because Linux cannot guarantee that kind of determinism.

BEFORE YOU COMMIT

- **Total system cost:** Compare the whole bill, not the chip. An MPU adds DDR memory, a power management IC, and usually a higher layer-count PCB, and it still often comes out ahead of the loaded MCU.
- **Boot and power behavior:** Decide whether the product can tolerate a Linux boot time and a proper shutdown sequence, or whether it needs to be alive within milliseconds of power-on.
- **Determinism:** List every task with a hard timing deadline. If any of them are in the microseconds, plan on an MCU for that task, possibly alongside the MPU.

The real mistake

The mistake is not choosing a bad microcontroller. It is choosing one that makes sense for a prototype but sabotages the product you actually want to build.

1.8 High-End Application Processors

TECHNOLOGY TO AVOID

WHY IT FAILS IN A PRODUCT

- These chips look incredible on spec sheets. They power flagship smartphones and tablets, with enormous performance, beautiful graphics, and every type of connectivity you can imagine.
- For a startup they are basically off-limits. The datasheets are locked behind NDAs, you usually will not get proper reference designs, and unless you are a huge customer you are not getting any meaningful support.
- Minimum order quantities can be tens of thousands of units, and integrating the software stack is a massive job even for experienced teams. I have seen entrepreneurs waste months trying to get basic documentation or battling buggy drivers that were never meant for the public.
- The same goes for dedicated AI and machine learning chips. They sound exciting, but they often come with high complexity and limited software support.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Stick with accessible platforms	Microcontrollers like the STM32 or ESP32 give you plenty of power for most products, with documentation you can download today.
Step up to a Raspberry Pi Compute Module	If you need more performance, a Compute Module gives you real horsepower with documentation and reference designs you can actually use.
Or an STM32MP1 or Allwinner T113	Both are Linux-capable processors with accessible documentation, and entry 1.7 covers when they make sense.
Prove any AI feature on mainstream hardware first	Get the machine learning feature working on a Compute Module or an ESP32-S3 with a mainstream framework before you consider a dedicated accelerator.

BEFORE YOU COMMIT

- **Datasheet test:** Try to download the full datasheet and reference manual without signing anything. If you cannot, you will not get the errata or the support either.
- **Reference design test:** Ask for a reference design with full schematics and a bill of materials, because a block diagram is not a reference design.
- **Minimum order:** Ask the minimum order quantity and lead time in writing before a single line is designed around the part.

The support test

Before you design in any processor, send the vendor a real technical question through their sales channel. If nobody answers a small customer before they have your business, nobody will answer after.

1.9

Custom IC Design

TECHNOLOGY TO AVOID

WHY IT FAILS IN A PRODUCT

- This one comes up more often than you would expect, especially with wearables or ultra-compact products. Founders get the idea that the only way to fit everything into a tiny enclosure is with a custom IC.
- In theory that is true, because custom silicon lets you combine radios, processors, and memory into one sleek package. In practice, custom chip design is a mountain you do not want to climb.
- You are talking about massive up-front costs, long lead times, highly specialized engineering, and extremely expensive mask sets, with a real chance of sinking your budget into a chip that never makes it past tape-out.
- Even when the chip works, you now own its supply chain, its testing, its qualification, and its errata, with none of the volume that makes those costs disappear for a big company.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Use off-the-shelf chips and proven modules	You can still build something compact, and you will not risk your budget on a chip that may never ship.
Look at system-in-package parts	Parts like Espressif's PICO series put the radio, the processor, the crystal, and the passives in one tiny footprint, and a chip-scale package like Nordic's WLCSP shrinks the chip itself, which together is most of the size benefit a custom chip would give you.
Get the size back with the PCB	Double-sided placement, smaller passives where they are justified, and a rigid-flex board shrink a product far more cheaply than custom silicon.

BEFORE YOU COMMIT

- **Do the layout exercise first:** Have your electrical engineer lay out the densest off-the-shelf version of the design before anyone says the word custom. The result is usually smaller than the founder assumed.
- **Budget and volume:** Unless you are building the next Apple Watch, have raised millions of dollars, and expect volumes in the millions of units, the math does not work.
- **Timeline:** Add years, not months, to the schedule for a custom chip, and ask whether the product still makes sense if it launches that late.

From someone who designed chips for a living

I spent over a decade designing chips at Texas Instruments, with a full team and a fab behind me, and a new chip still took years to get right. A startup does not have those years.

STAGE: PROTOTYPING

ARDUINO AND DEVELOPMENT BOARDS

If you are prototyping your product on one of these Arduino boards, be prepared to throw away months of firmware work when it is time to move to a custom PCB. They all work fine for early prototypes, which is exactly what makes them dangerous, because the board, and often the chip underneath, is not a foundation you can carry forward into production. None of them are designed to go inside a finished product either. Any FCC or CE certification they carry as dev boards does not transfer to your product, and an unshielded board with exposed headers can make emissions testing harder, not easier.

IN THIS PART

2.1	Arduino Uno R3, Mega 2560, and Nano	p. 15
2.2	Arduino Nano RP2040 Connect	p. 16
2.3	Arduino MKR Series	p. 17
2.4	Arduino Uno R4 WiFi	p. 18
2.5	Arduino Portenta H7	p. 19
2.6	Arduino Uno Q	p. 20

2.1 Arduino Uno R3, Mega 2560, and Nano

CLASSIC 8-BIT BOARDS

WHY IT FAILS AS A STARTING POINT

- The Uno is the board that got many of you started, which is exactly why it keeps showing up in early product designs where it does not belong. The ATmega328P at its heart is an old 8-bit chip that is underpowered, overpriced, and limited in I/O compared to modern 32-bit parts that cost less and do more.
- When people run out of pins on the Uno they reach for the Mega, built on the ATmega2560 with 54 digital I/O and more memory, but it shares the same old 8-bit architecture, so it fixes the pin count and nothing else.
- The Nano is physically smaller so it feels product-friendly, which is why people try to embed it directly inside their enclosures. It is still the same 328P with all the same limitations, just in a tinier package.
- People sometimes solder Nanos onto carrier boards inside their products, stacking PCBs and adding height, cost, and reliability problems like failure under vibration and thermal cycling. If you are soldering a dev board onto another board, you are ready for a custom PCB.

PROTOTYPE ON THIS INSTEAD

PROTOTYPE ON	WHY THIS BOARD
An Espressif ESP32 dev board	If your product needs Wi-Fi or Bluetooth, this is the board to start on, and the code moves cleanly to a pre-certified ESP32 module on your own PCB.
A Nordic nRF52840 DK	If you only need Bluetooth Low Energy, Nordic's development kit puts you on the chip family most BLE products ship with.
An STM32 Nucleo board or Raspberry Pi Pico	If the product needs no wireless, either one gets you prototyping on a production-capable chip for a few dollars.

BEFORE YOU COMMIT

- **Pick a chip with room to grow:** Choose a part based on the pins and memory you will need two revisions from now, not the exact count you need today, or you will be forced to switch chips or add an I/O expander later.
- **Keep the IDE if you want it:** The Arduino IDE works with almost every board in the table above, so moving off the Uno does not mean learning a new environment on day one.
- **Small-scale sales test only:** Embedding a Nano is acceptable for a handful of units to test demand, as long as you know a custom PCB comes next.

Why the manufacturer's boards win

Chip vendors' own dev boards are cheaper than Arduino boards, closer to what your production hardware will actually look like, and the code you write transfers much more cleanly when you move to a custom PCB.

2.2

Arduino Nano RP2040 Connect

TWO CHIPS FOR ONE JOB

WHY IT FAILS AS A STARTING POINT

- This board looks like a strong prototyping choice because the RP2040 inside it is actually a solid microcontroller, dual-core with plenty of I/O and under a dollar in production quantities. The problem is not the chip, it is the board wrapped around it.
- You are paying around 25 dollars for an RP2040 paired with a u-blox NINA module for Wi-Fi and Bluetooth, and that NINA module actually contains an ESP32 inside it, so you are stuck with two microcontrollers when a single ESP32 could have handled the whole job for a fraction of the cost.
- Two chips means two firmware images, two toolchains, and an interface between them that has to be debugged, and that architecture follows you onto the custom PCB if you do not change it now.
- In most cases, one microcontroller that does everything is going to be more cost-effective than paying for two separate chips, one for control and one for wireless.

PROTOTYPE ON THIS INSTEAD

PROTOTYPE ON	WHY THIS BOARD
An Espressif ESP32 dev board	One chip handles the application, the Wi-Fi, and the Bluetooth, and it maps directly onto a pre-certified module in production.
A Raspberry Pi Pico	If the wireless was never really needed, the Pico gives you the same RP2040 for a few dollars with no second chip in the way.
A Raspberry Pi Pico W	If you want to stay on the RP2040 and add Wi-Fi, the Pico W pairs it with an Infineon radio, but you are still carrying two chips, so weigh it against the ESP32.

BEFORE YOU COMMIT

- **Count the processors:** Draw the block diagram of the production design. If there are two programmable chips and one of them only does wireless, ask what the second one is for.
- **Price the production version:** Compare a single ESP32 module against an RP2040 plus a wireless module at your volume, including the second set of support parts.
- **Check the wireless certification path:** Confirm that whatever radio you prototype on has a pre-certified production module, so the prototype and the product share the same RF design.

The chip is fine, the board is the mistake

The RP2040 shows up as a recommended alternative elsewhere in this handbook. This entry is about paying for a board that bolts a second microcontroller onto it to do a job the ESP32 does alone.

2.3

Arduino MKR Series

IOT DEV BOARDS

WHY IT FAILS AS A STARTING POINT

- The MKR line is Arduino's attempt at IoT and product development, with Wi-Fi, LoRa, and cellular variants. They use the SAMD21, a 32-bit ARM Cortex-M0+ from Microchip, which is a fine microcontroller but nothing special at the 30 to 50 dollar price tag Arduino charges for these boards.
- The MKR Wi-Fi board pairs the SAMD21 with the same kind of u-blox NINA module covered in entry 2.2, which means you are once again paying for two microcontrollers when one ESP32 could handle the whole job for less.
- Where the MKR really falls apart is on the cellular side, because u-blox cellular modules are considerably more expensive than comparable options from Quectel or SIMCom, and cellular products usually have tight cost targets that make that premium hard to justify.
- The LoRa version is not much better, since the MKR WAN boards use a Murata module built around a Semtech SX1276, while the newer SX126x family is more flexible, more power-efficient, and better supported for new designs.

PROTOTYPE ON THIS INSTEAD

PROTOTYPE ON	WHY THIS BOARD
An Espressif ESP32 dev board	For the Wi-Fi products the MKR was aimed at, one ESP32 replaces the SAMD21 and the NINA module together.
A Quectel or SIMCom cellular evaluation kit	Prototype on the module family you can afford to ship, so the AT command set and the antenna design carry straight into production.
An SX126x-based LoRa module	Start on a module built around the current Semtech family, from a vendor that publishes regional certification, and see entry 3.7 before you pick one.

BEFORE YOU COMMIT

- **Price the module, not the board:** The number that matters is the cost of the wireless module you will ship, so get that quote before you prototype on anything.
- **Check the radio generation:** For LoRa, confirm the module is built on SX126x-class silicon, which includes the STM32WL. For cellular, confirm the technology is one your carriers will still support in five years, as covered in entry 3.8.
- **Plan the production board now:** Sketch the production block diagram before prototyping, so the dev board you choose actually resembles it.

An awkward middle ground

The MKR sits between two stools. It is too expensive for any real production volume and too limited for complex products, so it rarely ends up being the right answer for either.

2.4

Arduino Uno R4 WiFi

TWO CHIPS FOR ONE JOB

WHY IT FAILS AS A STARTING POINT

- The R4 WiFi is Arduino's newest mainline board and it looks like a big upgrade over the R3. It pairs a Renesas RA4M1 with an ESP32-S3 as a Wi-Fi and Bluetooth coprocessor, which means you are paying for two chips when one would do the job.
- The RA4M1 is a decent microcontroller, but it has a much smaller ecosystem than the STM32 or ESP32 families, with fewer reference designs, libraries, and community examples to pull from.
- The irony is that the coprocessor on the board is the more capable chip. The ESP32-S3 alone could run the whole product, so the RA4M1 adds cost and a second firmware image for peripherals, like 5 V I/O and CAN, that most products never use.
- You end up with a board that costs more than a bare ESP32-S3 dev board while adding a second chip with weaker long-term support.

PROTOTYPE ON THIS INSTEAD

PROTOTYPE ON	WHY THIS BOARD
An Espressif ESP32-S3 dev board	One chip covers the application, the Wi-Fi, and the Bluetooth, with a dual-core processor, native USB, and a large ecosystem.
The Arduino IDE to start, ESP-IDF to ship	The Arduino core works fine to get going, and as you get closer to production you will probably want the full ESP-IDF toolchain. The Arduino core can even run as a component inside an ESP-IDF project.
A pre-certified ESP32-S3 module for production	When you are ready to ship, move to a WROOM or MINI module with FCC and CE already done, and see entry 3.4 for why the shielded module matters.

BEFORE YOU COMMIT

- **One firmware image:** Confirm the production design has one programmable chip to update, secure, and support in the field, unless a second one earns its place with a peripheral the first cannot provide.
- **Ecosystem check:** Search for reference designs and libraries for the exact chip you will ship on. The ESP32-S3 has thousands, and that is the pool your engineers will draw from.
- **Certification path:** Confirm the S3 module you plan to ship carries modular FCC and CE approval, and design the prototype around the same module if you can.

The upgrade that is not one

The R4 WiFi looks like progress over the R3, but for a product it swaps one problem for another. The chip you actually want is already on the board, sitting there as the coprocessor.

2.5

Arduino Portenta H7

HIGH-END DEV BOARD

WHY IT FAILS AS A STARTING POINT

- The Portenta is Arduino's high-end board, built on the STM32H747 and aimed at industrial applications. These boards cost 100 to 120 dollars or more, which is a lot of money for a development board.
- The STM32H747 chip itself costs around 15 dollars in modest quantities, so you are paying roughly 7 to 8 times the chip price for the board, and even after you count the memory, radio, and PHY it carries, that is far more than a custom board costs at volume.
- Arduino markets the Portenta as production-ready, but that price makes it completely impractical at any real volume, and the ultra-high performance of the H7 is overkill for many of the products it ends up in.
- If your product truly needs STM32H7 performance, it is complex enough that you should not be relying on the Arduino software framework anyway.

PROTOTYPE ON THIS INSTEAD

PROTOTYPE ON	WHY THIS BOARD
An STM32H7 Nucleo or Discovery board	ST's own boards give you the same chip family for a fraction of the price, with the full peripheral set broken out.
STM32CubeIDE from day one	Building on ST's HAL and tools gives you real debugging, vendor drivers, and no framework layer to unwind later, which is what makes firmware production-ready, not the editor.
An STM32MP1 if you are pushing the H7	If you really need that level of performance, consider whether a microprocessor is a better fit than a high-end MCU that gets quite expensive, as covered in entry 1.7.

BEFORE YOU COMMIT

- **Do you need the H7 at all:** Write down the processing the product actually does. Most products that land on a Portenta would run comfortably on an STM32G4 or an ESP32-S3.
- **Price the bare chip:** Get a distributor quote for the H7 part you would ship and compare it with the board price, so the markup is a number rather than a feeling.
- **Toolchain decision:** Decide now whether the firmware will live in the Arduino framework or in STM32CubeIDE, because switching later means rewriting the driver layer.

Production-ready is a marketing term

A board can be reliable and well built and still be the wrong thing to ship, because at volume the markup on the packaging is what decides your margin.

2.6

Arduino Uno Q

QUALCOMM-BASED BOARD

WHY IT FAILS AS A STARTING POINT

- The Uno Q is Arduino's flashiest board yet, driven by their Qualcomm partnership. It combines a quad-core Qualcomm Dragonwing QRB2210 microprocessor with an STM32U585 microcontroller, runs Debian, supports camera and display connectivity, has AI acceleration, and starts around 50 dollars.
- That sounds impressive on paper, but the problem has nothing to do with the board and everything to do with Qualcomm, which has historically been one of the worst chip companies for startups and small teams to work with.
- They do not give you samples, they do not give you datasheets, they do not give you reference designs, and they do not give you meaningful support unless you are a multi-billion dollar customer buying millions of chips.
- If you do not need Linux or AI, the Uno Q is wildly overpowered and overpriced compared to prototyping on an ESP32 or STM32 dev board.

PROTOTYPE ON THIS INSTEAD

PROTOTYPE ON	WHY THIS BOARD
An ST STM32MP1 board	Linux, display, and camera support from a vendor with accessible documentation and an established supply chain.
An Allwinner T113 board	A low-cost Linux-capable processor with a large ecosystem of boards and reference designs you can get started on today.
A Raspberry Pi Compute Module 5	Built on the Broadcom BCM2712, but sold with full documentation, a design guide, and a supply chain built for small customers, which is the difference that matters.

BEFORE YOU COMMIT

- **Do you need Linux or AI at all:** If the product is basic embedded control, none of these boards belong in it, and an ESP32 or STM32 dev board is the right prototype.
- **The datasheet test:** Try to get the full processor datasheet and a reference design without signing an NDA. A distributor listing on Digi-Key does not change what happens after you ask for support.
- **Supply chain:** Ask who will sell you the processor in quantities of one thousand, at what lead time, and whether they will still be doing so in five years.

A listing is not a relationship

The Dragonwing has recently shown up on Digi-Key, which is a step in the right direction, but a distributor listing does not change decades of Qualcomm treating small customers like they do not exist.

STAGE: COMPONENT SELECTION

WIRELESS MODULES AND CONNECTIVITY

The wireless module you pick can kill your product before it ever ships. I am talking about FCC certifications that fall apart at the last minute, cellular networks that shut down mid-production, parts that go discontinued right when you are scaling, and clone modules that behave differently from the datasheet you designed around. Protocol selection is a bigger decision than module selection, so get the protocol right first, then pick the module to match, and then verify its certification yourself.

IN THIS PART

3.1	HC-05 and HC-06 Bluetooth Classic Modules	p. 22
3.2	HM-10 BLE Module	p. 23
3.3	Legacy BLE Modules (RN4020, BLE112, BLE113, CC2540)	p. 24
3.4	ESP-01 Wi-Fi Module	p. 25
3.5	Generic nRF24L01 Clone Modules	p. 26
3.6	Modules Whose Antenna Voids the Grant	p. 27
3.7	Generic Uncertified LoRa Modules	p. 28
3.8	SIM800 and SIM900 2G Cellular Modules	p. 29
3.9	Modules That Claim Pre-Certification but Are Not	p. 30
3.10	Custom RF Design, Especially Sub-GHz	p. 31
3.11	5G and Cutting-Edge Cellular Modules	p. 32

SECTION

PART 03

DOC NO.

NVR-01

ENTRIES

11

SOURCE

[PREDICTABLEDESIGNS.COM](https://predictabledesigns.com)

3.1 HC-05 and HC-06 Bluetooth Classic Modules

BLUETOOTH SPP

WHY IT FAILS IN A PRODUCT

- These are Bluetooth Classic modules built around the Serial Port Profile, or SPP, and they are the wrong choice for almost any consumer product that talks to a smartphone.
- iOS does not natively expose SPP to third-party apps, and the HC-05 and HC-06 cannot support Apple's MFi program at all, so there is no path to iOS compatibility with these modules. Half of your customers are locked out before you start.
- The market is flooded with counterfeits, so two modules with the same marking can behave completely differently in the field, and your firmware has no way to tell which one it is talking to.
- Most units sold have no legitimate FCC modular approval, which means you have no usable grant to inherit when it is time to certify your product.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
A Fanstel BLE module	My favorite vendor for Bluetooth products, building on Nordic nRF52 and nRF53 silicon with a verifiable FCC ID and datasheet-backed specs.
A Raytac or Laird BLE module	Both are solid options on modern Nordic silicon with real certification documentation you can look up yourself.
A Silicon Labs based module	If you prefer the Silicon Labs ecosystem, their BLE parts and the modules built on them carry the same kind of verifiable approvals.

BEFORE YOU COMMIT

- **Phone compatibility:** Confirm the product works with iOS and Android through BLE from day one, and drop any plan that depends on SPP.
- **FCC ID lookup:** Search the module's FCC ID in the FCC equipment authorization database and read the grant notes, as entry 3.9 explains.
- **Throughput reality check:** BLE is not Bluetooth Classic. If the product streams audio or bulk data, confirm the data rate you need is achievable over BLE before you commit.

Why I lean toward Nordic

Nordic's SDK, documentation, and community are the strongest in BLE, and modules built on their chips give you a well-trodden path from prototype to a certified product.

3.2

HM-10 BLE Module

BLE 4.0 ERA

WHY IT FAILS IN A PRODUCT

- The HM-10 is a BLE 4.0 era part with serious clone problems lurking behind it. The chip inside may or may not be the TI CC2540 or CC2541 the original design used, because the market is full of clones running completely different silicon.
- The AT command set everyone designs around is proprietary to the original manufacturer's firmware, and clones do not always replicate it correctly.
- Picture this: you design around a specific set of AT commands, you scale to production, and your second batch behaves differently because your assembler sourced from a different clone supplier. That problem shows up after you have already shipped.
- BLE 4.0 also lacks the features newer phones and products expect, like the longer range and higher data rate modes that arrived with BLE 5, so the module limits the product even when it works.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
A Fanstel BLE module	My go-to recommendation, built on modern Nordic nRF52 silicon with real certification and a supply chain you can trace.
A Raytac or Laird BLE module	Solid choices on the same Nordic silicon, with published FCC IDs and datasheets that match the part you receive.
Nordic's own SDK instead of AT commands	Running your application on the module's processor removes the AT command layer entirely, though it does not remove the need for a traceable source.

BEFORE YOU COMMIT

- **Silicon check:** Ask the vendor which chip is inside the module and confirm it against the datasheet. If they cannot say, that is your answer.
- **Lock the source:** Specify the exact manufacturer part number on your bill of materials and mark it as no substitutions allowed, so an assembler cannot swap in a clone.
- **BLE features, not just the version:** Confirm the module supports the specific features the product needs, like the 2M or Coded PHY, secure pairing, and your throughput, since a version number alone guarantees none of them.

The second batch problem

Clone modules are at their most dangerous on the reorder, because the first batch works, the product ships, and the batch that behaves differently is the one already in customers' hands.

3.3

Legacy BLE Modules (RN4020, BLE112, BLE113, CC2540)

END-OF-LIFE MODULES

WHY IT FAILS IN A PRODUCT

- I used to love a lot of these modules and designed with them frequently, but that was quite a while ago, and they are all outdated now with better alternatives available.
- They are BLE 4.0 or 4.1 era parts missing features that newer products often need, and many are marked not recommended for new designs or flat out discontinued, with SDK support winding down.
- You design in something like the RN4020 because a tutorial used it or it was in an old bill of materials, and two years later you cannot source the part and your firmware is tied to that module's stack.
- An old module also means an old security profile, and connected product rules like the ones in entry 11.9 increasingly assume current pairing and encryption features.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
A current BLE 5.x module on Nordic nRF52 or nRF53 silicon	These carry the features, the security, and the SDK support a product launched today will need for its whole life.
A module from a vendor with a published lifecycle policy	Fanstel, Raytac, Laird, and u-blox publish lifecycle status, which is what lets you plan around a part instead of being surprised by it.
Silicon Labs or Espressif for other ecosystems	If the rest of the product is already on Silicon Labs or ESP32, their current BLE 5 parts keep the design inside one toolchain.

BEFORE YOU COMMIT

- **Silicon inside:** Check what chip is actually inside any BLE module before you design it in, then check that chip's lifecycle status on the silicon vendor's site.
- **Spec version and SDK:** Confirm which BLE version the module supports and look at how recently the SDK was updated. A repository that went quiet two years ago is a warning.
- **Exit plan:** Pick a module whose vendor offers a pin-compatible successor, so an end-of-life notice means a part swap rather than a redesign.

Tutorials do not track lifecycles

A tutorial written five years ago is still a tutorial today, and the module it used is often not a module anymore. Always check the part before you check the code.

3.4

ESP-01 Wi-Fi Module

UNSHIELDED MODULE

WHY IT FAILS IN A PRODUCT

- The ESP-01 is a Wi-Fi module built on the ESP8266 and it is everywhere in tutorials and maker projects. It is great for learning Wi-Fi basics and a terrible choice for a real product.
- GPIO is minimal and the antenna is a poorly tuned trace, but the real problem is that the entire form factor lacks an RF shield.
- An RF shield is a requirement for full FCC modular approval, so no standard ESP-01, whether genuine or clone, can hold a full modular grant. That leaves you paying for intentional radiator testing on your own product.
- The 8266 underneath it has all the limitations covered in entry 1.1, so even a shielded version would be a step backward from current ESP32 silicon.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
ESP32-C3-WROOM-02 or ESP32-C3-MINI-1	The cheapest properly shielded modules on modern silicon, with FCC modular approval and well-documented production use.
ESP32-C6-WROOM-1 or MINI-1	Adds Wi-Fi 6, Thread, and Zigbee for products that need a mesh or a newer router standard.
ESP32-S3-WROOM-1 or MINI-1	The choice when the product needs more processing, native USB, or a display alongside the radio.

BEFORE YOU COMMIT

- **Shield check:** If you can see the chip and the antenna trace with no metal can over the circuitry, the module cannot hold a full modular grant, whatever the listing claims, and any limited grant carries host restrictions you have to read.
- **Grant lookup:** Find the module's FCC ID on the module label and confirm it in the FCC database, as entry 3.9 describes.
- **Antenna variant:** Choose the PCB antenna version unless you have RF expertise, and follow the keepout in the datasheet exactly, as covered in entry 10.10.

The 8266 is not going away, but you should

Espressif's longevity commitment for the 8266 runs until 2029, so the chip is not the emergency. The unshielded ESP-01 form factor is, because it can never carry a full modular grant.

3.5 Generic nRF24L01 Clone Modules

PROPRIETARY 2.4 GHZ

WHY IT FAILS IN A PRODUCT

- I mean the cheap clones flooding the market, not legitimate Nordic silicon used properly. These modules use a proprietary 2.4 GHz protocol, not BLE, not Zigbee, not Thread, which means you are building an island that cannot talk to anything else.
- The clones have no inheritable approvals or trustworthy compliance documentation, the range claims on the listings are basically fiction, and there is no standardized ecosystem for interoperability.
- If you are reaching for an nRF24L01, it usually means you have not decided on a real protocol yet, and protocol selection is actually a bigger decision than module selection.
- A proprietary link also means you write and maintain the whole protocol stack yourself, including the pairing, the security, and the error handling that a standard already solved.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Pick the protocol before the part	Decide whether the product needs to talk to a phone, a hub, a cloud, or only to its own peers, and let that answer choose between BLE, Wi-Fi, Thread, or Zigbee.
Use plain BLE for point to point	For most point-to-point applications, a BLE module on Nordic silicon gives you a standard link, phone compatibility, and a certified module.
Use BLE mesh, Thread, or Zigbee for small meshes	These are standardized, certified, and supported by current silicon, and the ESP32-C6 and Nordic nRF52840 both cover them.

BEFORE YOU COMMIT

- **Who does the product talk to:** List every device that must connect to it over its life, including the phone app you have not built yet, and confirm the protocol reaches all of them.
- **Approvals:** Confirm the module you pick carries a real modular grant, since a proprietary clone module never does.
- **Range, measured:** Ignore the listing and test the range in the actual enclosure at the actual data rate before you commit.

An island is expensive

Every proprietary link you build is one more thing you alone have to secure, certify, and maintain for the life of the product, and none of your customers' other devices will ever join it.

3.6

Modules Whose Antenna Voids the Grant

RF PIN AND CONNECTOR MODULES

WHY IT FAILS IN A PRODUCT

- This is a whole category where the antenna and certification story falls apart in practice. It includes raw Semtech SX1276 or SX1262 breakouts and modules with an RF trace pin, where the modular grant only stays valid if you replicate the vendor's exact microstrip layout and stackup, which beginners almost always deviate from.
- It also includes modules with a U.FL or SMA connector where the FCC grant lists specific antenna types and maximum gains, which integrators either ignore or do not bother to check before picking a completely different antenna.
- The moment you deviate from a module's grant antenna conditions, the modular approval no longer covers your integration. That can mean full intentional radiator testing for the US alone at roughly 8 to 15 thousand dollars per radio, with multi-region compliance pushing past 30 thousand.
- Antenna matching on RF trace pin modules usually requires a vector network analyzer and RF expertise most product creators do not have, plus impedance-controlled traces and proper antenna placement.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
A module with an integrated PCB or chip antenna	The antenna is part of the certified module, so as long as you follow the datasheet keepout and the other grant conditions, the grant covers the radio in your product.
A connector module used exactly as the grant allows	If you must use an external antenna, match the antenna type listed in the grant and stay at or below the listed gain, with no exceptions.
A vendor reference layout, copied exactly	If an RF pin module is unavoidable, copy the vendor's microstrip layout and stackup line for line, and budget for pre-compliance testing.

BEFORE YOU COMMIT

- **Read the grant, not the listing:** Pull the grant from the FCC database and read the antenna conditions. If the antenna you plan to use is not listed by type and gain, you are buying an RF project.
- **Keepout and placement:** Confirm the enclosure, the battery, and every cable stay out of the antenna keepout, since a wire running past the antenna can detune it just like copper does.
- **Budget the alternative:** Price full intentional radiator testing for every market you sell in, so the saving on a cheaper module is judged against a real number.

The simple rule

If a module does not ship with an integrated antenna or a grant listing the antenna type and gain you intend to use, you are buying yourself an RF project.

3.7 Generic Uncertified LoRa Modules

SUB-GHZ LONG RANGE

WHY IT FAILS IN A PRODUCT

- LoRa is useful for long range, low bandwidth applications, but the frequency bands vary by region, and cheap LoRa modules often have no legitimate regional certification.
- Even when the hardware works fine across frequency bands, the firmware, channel plan, and duty cycle requirements are region-specific compliance gates that no amount of hardware matching can fix.
- You design a LoRa product, ship into a region where the module does not meet local compliance requirements, and your product gets pulled.
- Sub-GHz is also where pre-certified options are thinnest, so the temptation to roll your own radio is strongest exactly where the RF work is hardest, as entry 3.10 explains.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
A Murata LoRa module	Murata builds certified modules on Semtech silicon with a real datasheet and regional variants you can specify by part number.
A RAK Wireless module	RAK offers certified modules and a large LoRaWAN community, which shortens the firmware side considerably.
An STM32WL-based module such as the RAK3172	ST's STM32WL puts a Semtech-class sub-GHz radio inside the microcontroller, and modules built on it come in explicit regional variants with certification.

BEFORE YOU COMMIT

- **Regional variant on the BOM:** Put the exact regional part number on the bill of materials for each market, since the 915 MHz and 868 MHz versions are usually different orderable parts with different grants.
- **Channel plan and duty cycle:** Confirm the firmware enforces the regional channel plan and duty cycle limits, because the module hardware will happily transmit outside them.
- **Silicon generation:** Confirm the module is built on SX126x-class silicon, which includes the STM32WL, rather than the older SX127x, for the power and flexibility reasons covered in entry 2.3.

Compliance lives in the firmware too

With sub-GHz radios, a certified module gets you the hardware half. The channel plan, the duty cycle, and the output power settings in your firmware are the other half, and the test lab checks both.

3.8 SIM800 and SIM900 2G Cellular Modules

2G CELLULAR

WHY IT FAILS IN A PRODUCT

- 2G networks are shut down or shutting down in most major markets, so any product built on these is basically obsolete the day it ships.
- Ship a SIM800 to the US and you have no dependable carrier path. Ship to Australia and it flat out cannot connect. Ship to Europe and coverage is fragmented and shrinking.
- These modules are also power-hungry compared to modern low-power cellular, with transmit bursts that collapse a small battery, which is a problem the LTE-M and NB-IoT standards were designed to fix.
- A cellular product usually has a service contract and a multi-year life attached to it, so a radio technology on a shutdown schedule is a liability you are signing your customers up for.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
An LTE-M or NB-IoT module	For low-data IoT, modules from Quectel, u-blox, or Telit give you a low-power modern radio, but check regional coverage carefully because deployment is still uneven.
An LTE Cat-1 bis module	Often the better default for global coverage, since it runs on standard LTE networks wherever they exist, and all three of those vendors offer modules that support it. You still pick bands and carrier variants per market.
A module certified for your carriers	Whatever you pick, make sure it is certified for your target regions and carriers, and see entry 3.11 before you consider 5G.

BEFORE YOU COMMIT

- **Carrier coverage map:** Check LTE-M and NB-IoT coverage for every country you will sell in, on the carrier's own site, not the module vendor's.
- **Carrier certification:** Ask the module vendor for the list of carriers the module is certified with, and confirm your product does not need its own carrier certification on top.
- **Shutdown dates:** Search for the published sunset dates of every network technology the module depends on, and compare them with the product's expected life.

Obsolete on day one

A 2G module can pass every test on your bench, connect on a shrinking network for a while, and then turn every unit you sold into a brick on a date the carrier chose.

3.9

Modules That Claim Pre-Certification but Are Not

CERTIFICATION

WHY IT FAILS IN A PRODUCT

- This is the most important entry in the part, because the lesson applies to every wireless decision you will ever make. These are listings claiming FCC certified or CE certified right in the title, but the paperwork does not hold up.
- The FCC IDs may be fake, belong to a completely different module, or have grant conditions that do not actually cover your integration.
- So you design around a module that looks pre-certified, build a thousand units, and at final compliance your lab catches it. Now you are either paying for full intentional radiator testing at 8 to 15 thousand dollars per radio, or scrapping inventory.
- CE is self-declared, so a CE mark on a listing means even less than an FCC claim. What matters is whether the manufacturer can hand you the test reports behind it.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Look up the FCC ID yourself	Search the FCC equipment authorization database for the ID printed on the module, and confirm the grantee and the product description match the module in your hand.
Read the grant notes and conditions	Pay attention to whether it is a full modular approval or a limited one, because limited approvals come with host-specific conditions that may not cover your design.
Buy from manufacturers who publish the documents	If you cannot find a real FCC ID in the database and a test report from the manufacturer, it does not count as certified.

BEFORE YOU COMMIT

- **Database match:** The FCC ID, the grantee name, and the module photos in the database all match the part you received.
- **Grant type:** The grant says single modular or limited modular, and you have read every condition that a limited grant imposes on the host product.
- **Antenna and power:** Your antenna and output power settings sit inside what the grant lists, as entry 3.6 covers.

The rule is simple

If you cannot find a real FCC ID in the FCC equipment authorization database, it does not count as certified, no matter what the listing title says.

3.10 Custom RF Design, Especially Sub-GHz

TECHNOLOGY TO AVOID

WHY IT FAILS IN A PRODUCT

- Trying to do your own RF design, especially in the sub-gigahertz range, is a huge gamble. At 2.4 GHz you can at least find pre-certified modules for Wi-Fi, Bluetooth, and Zigbee, but once you drop into 433 MHz or 900 MHz territory, you are usually on your own.
- That means designing antennas from scratch, tuning matching networks, and trying to pass FCC or CE testing without much vendor support.
- RF is tricky. Your layout, your enclosure, even the battery wire running too close to the antenna, all of it can break performance, and each of those changes can send you back to the test lab.
- Without a vector network analyzer, a spectrum analyzer, and someone who knows how to use them, you cannot even see the problem you are trying to fix.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Use a certified module	Unless you have serious RF expertise and test equipment, buy the radio as a module with an integrated antenna and a modular grant, and follow its keepout exactly.
Choose technologies with reference designs	Stick to wireless standards where the silicon vendor provides proven reference designs and antenna solutions you can actually use.
Bring in RF help for anything custom	If the product truly requires a custom radio, hire an RF engineer for the antenna and matching work and budget for pre-compliance scans before the real test.

BEFORE YOU COMMIT

- **Is there a module for it:** Before you design any radio, search for a certified module in the band you need. If one exists, that is the design.
- **Enclosure and cables in the model:** Test the antenna inside the real enclosure with the real battery and cables in place, not on a bare board.
- **Pre-compliance budget:** Put a pre-compliance scan in the schedule and the budget before the first revision, so an EMI surprise costs days rather than a launch.

The battery wire that broke the product

One of the most common RF failures I see is not the antenna at all. It is a wire or a metal bracket routed through the keepout after the RF work was done, which detunes everything.

3.11 5G and Cutting-Edge Cellular Modules

TECHNOLOGY TO AVOID

WHY IT FAILS IN A PRODUCT

- 5G sounds like the future, and for some applications it is. But most startups do not need it, and trying to integrate it usually ends up being a huge drain on resources.
- These modules are expensive. They often require specialized SIM provisioning, thermal design considerations, and compliance testing with cellular carriers.
- Just because a module is pre-certified does not mean your end product will sail through. If your antenna tuning, ground plane, or enclosure affects performance, you could be forced to go through carrier testing all over again.
- The power and thermal budget of a 5G radio is far beyond what a small battery-powered product can support, so the technology decides the enclosure and the battery before you have designed either.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
A pre-certified LTE Cat-M1 or NB-IoT module	Cheaper, easier to source, and far more forgiving for low-to-medium bandwidth use cases, with low-power modes built for battery products.
A pre-certified 4G LTE Cat-1 bis or Cat-4 module	When you need real throughput, standard LTE gives you the bandwidth with a mature certification path and global coverage.
5G only when the product truly requires it	Real-time video or high-throughput data streaming are the cases where 5G earns its cost, and even then plan on carrier testing.

BEFORE YOU COMMIT

- **Throughput, in numbers:** Write down the sustained data rate the product actually needs. If it is under a few megabits per second, LTE covers it.
- **Carrier testing:** Ask the module vendor and your target carrier what testing the end product needs, and get the cost in writing.
- **Thermal and power:** Model the module's peak transmit power against the enclosure and the battery, because those numbers decide whether the product can be sealed and how long it runs.

More pain than it is worth

Unless your product truly requires high throughput and real-time video or data streaming, 5G is usually more pain than it is worth, and the customer cannot tell the difference.

STAGE: COMPONENT SELECTION

SENSORS YOU SHOULD NEVER USE

If you learned electronics through beginner kits and tutorials, at least one of these sensors is probably sitting in your prototype right now. There is nothing wrong with them for learning. But once you make the jump from a bench prototype to a product you can manufacture, certify, and sell, every sensor gets judged on accuracy, lifecycle, compliance, power, and safety, and every one of these fails at least one of those filters. Each entry gives you a production-grade replacement, usually for close to the same cost in volume.

IN THIS PART

4.1	DHT11 and DHT22	p. 34
4.2	HC-SR04 Ultrasonic Module	p. 35
4.3	HC-SR501 PIR Motion Module	p. 36
4.4	MQ-Series Gas Sensors (MQ-2, MQ-135)	p. 37
4.5	MPU-6050 IMU	p. 38
4.6	Cadmium Sulfide Photoresistors (LDRs)	p. 39
4.7	ACS712 Current Sensor Modules on Mains Voltage	p. 40

4.1 DHT11 and DHT22

TEMPERATURE AND HUMIDITY

WHY IT FAILS IN A PRODUCT

- The DHT11 is the little blue temperature and humidity sensor in basically every starter kit ever sold. It costs about a dollar, there are ten thousand tutorials for it, and it was never designed to go inside a real product.
- Its accuracy is only rated to plus or minus 2 degrees C, the humidity readings are even worse, and the unit-to-unit variance is wide enough that two sensors from the same batch will noticeably disagree.
- It talks over a nonstandard one-wire protocol with strict timing requirements, so readings randomly get corrupted, and a bit-banged driver that gets interrupted mid-pulse makes it worse, which is exactly what happens under an RTOS.
- These are cranked out by countless factories of unknown origin, so you never really know what you are getting, and the DHT22 shares the same flaky protocol, the same variance problem, and the same questionable supply chain.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
Sensirion SHT40 or SHT41	Tiny, factory-calibrated sensors that communicate over standard I2C, with accuracy specified across the full operating range rather than at one comfortable point.
TI HDC3020	Another fully specified option from a major manufacturer, with a similar I2C interface and a proper datasheet.
Bosch BME280	If you want barometric pressure sensing in the same package, the BME280 adds it, but buy it through an authorized channel for the reason in entry 13.2.

BEFORE YOU COMMIT

- **Accuracy over the range:** Read the accuracy table at the temperatures and humidities the product will actually see, not just the headline number at 25 degrees C.
- **Placement:** Keep the sensor away from the processor, the regulator, and the battery, and give it an airflow path, or the most accurate part on the board will report the board's temperature.
- **Source:** Buy from an authorized distributor and put the exact part number on the bill of materials with no substitutions, since counterfeit modules are common.

Cost check

In production volumes you are talking a couple of dollars, which is a small price to pay for temperature and humidity data you can trust.

4.2

HC-SR04 Ultrasonic Module

DISTANCE SENSING

WHY IT FAILS IN A PRODUCT

- The HC-SR04 is the classic ultrasonic module with two silver cans that look like a pair of eyes. It is a 5 V module living in a 3.3 V world, so you need level shifting just to connect it safely to most modern microcontrollers.
- It is a large through-hole module rather than a component you can place in your own PCB layout, and it eats a ridiculous amount of space inside an enclosure.
- There is no temperature compensation, and since the speed of sound changes with temperature, your distance readings drift with the weather. Condensation outdoors and the fragile exposed transducers make it a sensor that only behaves in a climate-controlled room.
- Once you put it behind an enclosure or a grille, the acoustics change and your readings shift all over again, so the prototype numbers never survive the industrial design.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
ST VL53L1X or VL53L4CD	Time-of-flight sensors measure how long it takes light to bounce back from a target, and they are excellent for short-range distance sensing in a solderable package.
MaxBotix ultrasonic sensors	If you truly need ultrasonic, these product-grade sensors are not necessarily cheap, but their spec sheet and reliability are exactly what you are paying for.
A reflective IR sensor	If all you need is basic presence detection, a simple reflective IR sensor does the job for pennies.

BEFORE YOU COMMIT

- **Range and target:** Check the datasheet range against the surface you are measuring, since dark or shiny targets shorten the range of an optical sensor.
- **Cover window:** If the sensor sits behind a window in the enclosure, follow the vendor's cover glass guidance for thickness, air gap, and crosstalk, and test with the real window.
- **Ambient light:** Test optical sensors in direct sunlight if the product will ever see it, because that is where the cheap ones fall apart.

Ultrasonic is fine, this module is not

Ultrasonic sensing is a perfectly valid approach for products. The problem is this specific two-dollar module and everything about how it was built.

4.3 HC-SR501 PIR Motion Module

MOTION DETECTION

WHY IT FAILS IN A PRODUCT

- The HC-SR501 is the white-domed module that gives you working motion detection in an afternoon for about two dollars, and that convenience is exactly what makes it a trap.
- The sensitivity and timing are set with two trimpots, which are tiny screw-adjusted dials, and there is no way to replicate those settings consistently across thousands of units in production.
- The retrigger behavior is inconsistent, the onboard regulator has its own quirks, and false triggers are a constant complaint, especially with a Wi-Fi antenna nearby.
- In a security or lighting product, every false trigger chips away at your customer's trust until they stop using the product completely.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
Panasonic EKMC series	A clean digital output, specified sensitivity, a whole range of lens options, and a package designed to be soldered directly into products.
Panasonic EKMB series	Made specifically for battery products, with current draw of just a few microamps.
ST STHS34PF80	If you want true presence detection, where the sensor still knows a person is there after they stop moving, this infrared sensor does it over I2C.

BEFORE YOU COMMIT

- **Lens and field of view:** Choose the lens for the detection pattern the product needs, since the lens decides the range and the zones, not the firmware.
- **Radio separation:** Keep the sensor and its wiring away from the Wi-Fi or cellular antenna, and test for false triggers during transmit bursts.
- **Timing in firmware, sensitivity in the part:** Hold time and retrigger logic live in firmware where they are identical on every unit. A digital-output PIR has its detection sensitivity set by the sensor and lens you choose, so pick those for the range you need, and never rely on a trimpot for either.

No trimpots in production

Any setting that lives on a screw-adjusted dial is a setting that will be different on every unit, and one that nobody can adjust once the enclosure is closed.

4.4 MQ-Series Gas Sensors (MQ-2, MQ-135)

GAS AND AIR QUALITY

WHY IT FAILS IN A PRODUCT

- The first dealbreaker is power, because the heater inside these sensors burns roughly 750 to 900 mW continuously. A battery-powered product is dead on arrival, since this one sensor drains more power than the entire rest of your design combined.
- They need a burn-in period of a day or more before readings settle, they drift constantly over their lifetime, and they respond to practically every gas at once.
- Without calibration against reference gases, which you cannot realistically do, the numbers coming out of these sensors mean almost nothing.
- You can never use these for any kind of safety function. Gas and carbon monoxide detection is a certified category with strict safety standards, and if your product claims to protect people, that liability lands directly on you.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
Sensirion SGP40 or SGP41	For general air quality sensing, these measure volatile organic compounds, the airborne chemicals given off by paints and cleaners. Sensirion's algorithm turns the raw signal into a VOC index, and the SGP41 adds a NOx channel.
Sensirion SCD40 or SCD41	If you want real carbon dioxide measurement, these give you it in a tiny package with a proper accuracy specification.
SPEC Sensors or Figaro electrochemical cells	For detecting one specific gas, an electrochemical cell is the way to do it, with a defined target gas and a defined lifetime.
A certified detection module	It costs more, but it comes with the testing and documentation that regulators, retailers, and insurance companies demand for any safety claim.

BEFORE YOU COMMIT

- **Safety claim or not:** Decide in writing whether the product makes any safety claim. If it does, the detector has to be assessed to the applicable safety standard, which usually means a certified module.
- **Sensor lifetime:** Electrochemical cells have a shelf life and an operating life. Confirm both against the product's expected life and plan the replacement or the end of life.
- **Power budget:** Add the sensor's heater or duty-cycled current to the power budget before anything else, since it usually dominates.

Never a safety function

If your product claims to detect smoke, gas, or carbon monoxide to protect a person, it has to be a certified detector. There is no version of an MQ sensor that gets you there.

4.5

MPU-6050 IMU

MOTION TRACKING

WHY IT FAILS IN A PRODUCT

- The MPU-6050 is different from everything else on this list, because it works pretty well. It combines an accelerometer and a gyroscope in a single part, with a decade of tutorials and mature software libraries behind it.
- TDK InvenSense officially flagged the 6050 as not recommended for new designs years ago. The part is discontinued, and a lot of what is sold on hobby modules today is remarked or flat-out counterfeit silicon.
- Counterfeit silicon means every new reel you buy can behave a little differently, which turns your firmware tuning into a guessing game.
- You cannot build a production product around a discontinued part, because you are one supply hiccup away from a stopped production line.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
TDK ICM-42688-P	The modern successor from the same company that made the MPU-6050, with better specs and lower power.
Bosch BMI270	Optimized for low power, which makes it a strong choice for wearables and battery products.
ST LSM6 family (such as the LSM6DSO)	A wide range of options with mature software support, and all of them are actually in production.

BEFORE YOU COMMIT

- **Lifecycle status:** Check the manufacturer's product page for the exact part number. Not recommended for new designs means the clock is already running.
- **Interface voltage:** Confirm the I/O voltage of the replacement, since some modern IMUs run their data lines at 1.8 V and need level shifting to talk to a 3.3 V microcontroller.
- **Firmware abstraction:** Put the IMU behind a driver interface in the firmware, so the next part swap is a driver change rather than a rewrite.

A warning from the real world

One Hardware Academy member built his design around the 9-axis MPU-9250 and had to switch parts mid-development when it went end of life. The swap was not trivial, because the replacement ran its data lines at 1.8 V, so he had to add level shifting just to talk to the new part.

4.6

Cadmium Sulfide Photoresistors (LDRs)

AMBIENT LIGHT

WHY IT FAILS IN A PRODUCT

- Photoresistors, also called LDRs, are about as simple as a sensor gets, just a two-lead part whose resistance changes with the light hitting it. The only reason they are still everywhere is inertia.
- The classic photoresistor is made from cadmium sulfide, and cadmium is restricted under RoHS, the European regulation that limits hazardous substances in electronics. In practice these photoresistors cannot go into a product sold in the EU.
- Even if compliance were not a problem, they are terrible sensors anyway. They respond slowly, nobody fully specifies which wavelengths they react to, and the part-to-part variance is so wide they barely qualify as sensors.
- A resistive sensor also needs an ADC channel and a divider that drifts with temperature, so the reading you get is never in units anyone can act on.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
Vishay VEML7700	A popular first choice with a wide dynamic range and calibrated readings in lux over I2C.
TI OPT3001	Designed to match the spectral response of the human eye, which matters for anything that adjusts brightness for a person.
Lite-On LTR-303	A compact, low-cost option over I2C, with two-channel readings you convert to lux in firmware using the datasheet formula.

BEFORE YOU COMMIT

- **RoHS declaration:** Get the RoHS declaration for every sensor in the product, not just the light sensor, since cadmium is not the only restricted substance.
- **Window transmission:** If the sensor sits behind tinted plastic, measure how much light gets through and confirm the sensor still reads correctly at the darkest level the product must detect.
- **Field of view:** Check the sensor's angular response against where light will actually come from in the enclosure.

Cost check

All three give you calibrated light readings over I2C that convert to lux with the datasheet formula, they are fully RoHS compliant, and in volume they cost around a dollar or less.

4.7

ACS712 Current Sensor Modules on Mains Voltage

AC CURRENT SENSING

WHY IT FAILS IN A PRODUCT

- The ACS712 module is the cheap measure-AC-current board at the heart of every DIY energy monitor tutorial. The Allegro chip itself is a legitimate part. The hobby breakout board it sits on is not.
- These modules route mains voltage through PCB traces with nowhere near the creepage and clearance that safety standards require, which is the minimum spacing between high voltage and everything else on the board.
- There is no reinforced isolation, and the screw terminals are often completely exposed. Inside a product, that is a shock hazard, a fire hazard, and an automatic certification failure all rolled into one.
- This is the one category where a bad sensor choice can actually hurt somebody, so treat anything that touches mains voltage with the respect it deserves.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
Shunt resistor with TI AMC1302	A shunt paired with an isolated amplifier, laid out on your own board with the proper spacing, is the standard approach.
A current transformer	The better choice when you would rather keep mains off your board entirely, since it clamps around a wire and gives you an isolated measurement by design, as long as its secondary always has its burden resistor connected.
Allegro ACS37800	A newer part made to be integrated correctly into your own layout, with power monitoring built in.
A certified energy metering module	A certified module is the safe option when you would rather buy the safety engineering instead of designing it.

BEFORE YOU COMMIT

- **Creepage and clearance:** Look up the required spacing for your mains voltage and pollution degree in the safety standard for your product category, and design the layout to it before routing anything.
- **Isolation rating:** Confirm the isolated amplifier or transformer is rated for reinforced isolation at your working voltage, with the certificate to prove it.
- **Do you need to measure mains at all:** A certified external supply keeps the power supply out of the enclosure, but a current measurement on the mains side still needs its own isolation design.

Safety warning

Treat anything that touches mains voltage with the respect it deserves, because this is the one category where a bad sensor choice can actually hurt somebody.

STAGE: POWER ARCHITECTURE

BATTERIES AND CHARGING

Every battery in this part works perfectly on your bench, and every single one of them has ruined somebody's product. One of them will get your shipment refused by your freight company before it ever leaves the country. Another puts your product under a federal safety law most people building products have never heard of. And unlike most design mistakes, a bad battery choice does not just mean a product that will not work. It can physically hurt your customer, and your company is the one liable when it does.

IN THIS PART

5.1	9V Alkaline	p. 42
5.2	Alkalines in High-Drain Products	p. 43
5.3	Coin Cells for Anything Beyond Tiny Loads	p. 44
5.4	Sealed-In Batteries	p. 45
5.5	Unprotected 18650 Cells	p. 46
5.6	Bare LiPo Pouch Cells	p. 47
5.7	Uncertified Lithium Packs	p. 48
5.8	Custom Battery Packs and Chemistries	p. 49
5.9	TP4056 Charging and Protection Boards	p. 50
5.10	Wireless Charging (Qi and Others)	p. 51

SECTION

PART 05

DOC NO.

NVR-01

ENTRIES

10

SOURCE

[PREDICTABLEDESIGNS.COM](https://predictabledesigns.com)

5.1

9V Alkaline

PRIMARY CELL

WHY IT FAILS IN A PRODUCT

- The 9V alkaline is the battery almost everybody reaches for on their first prototype, and it is the first one I would get rid of. It looks efficient, but it only holds around 500 to 600 mAh, and once you account for the voltage difference you are paying three to four times more per unit of energy than with a couple of AA cells.
- Your electronics almost certainly run on 3.3 V or 5 V. Drop 9 V down to 3.3 V with a simple linear regulator and you throw away about two thirds of your energy as heat. A switching regulator fixes the efficiency, but now you have added cost, parts, and noise to solve a problem you created by picking the wrong battery.
- The snap connector is a mechanical weak point, since those thin wires flex every time your customer changes the battery, and eventually one breaks off inside the enclosure.
- A 9V is really six tiny cells in series, which is why the capacity is so low and why the internal resistance climbs so fast under load.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
Two AA cells in series	Around 3 V, which many microcontrollers and sensors run on directly with no regulator at all, from cells your customer can buy anywhere.
A small regulator only if you need one	Dropping 3 V to 1.8 V with a cheap linear part, or boosting it up to 3.3 V, is a much smaller job than stepping down from 9 V.
A single-cell lithium polymer with a charger IC	If you would rather go rechargeable, this is a better fit for almost any modern product, with the protection covered in entry 5.6.

BEFORE YOU COMMIT

- **Energy per dollar:** Compare watt-hours per dollar across the candidate batteries, not just the printed voltage, and the 9V loses every time.
- **Operating voltage range:** Confirm every chip in the design runs from the two-cell range, from 3.2 V fresh down to the cutoff you choose, before you drop the regulator.
- **Contacts and holder:** Pick a battery holder with proper spring contacts rated for the insertion cycles the product will see, since the holder is now the mechanical weak point.

Cost check

Two AA cells hold more usable energy than a 9V at a fraction of the price, especially under load, and they remove a regulator from the bill of materials in many designs.

5.2 Alkalines in High-Drain Products

PRIMARY CELL

WHY IT FAILS IN A PRODUCT

- Alkaline AA and AAA cells belong in wall clocks and TV remotes, not in any product that pulls real current. The AA format itself is fine. It is the alkaline chemistry that gives out when your product leans on it.
- The problem is internal resistance, which eats into the voltage the second your product starts drawing power. A fresh alkaline AA sits at 1.5 V doing nothing, then sags toward 1.1 V under a heavy load, and that internal resistance keeps climbing as the cell drains. Your product browns out and resets while there is still plenty of capacity inside the battery.
- Cold weather makes it worse, since alkaline performance drops off badly near freezing, which is why I steer anyone building outdoor products away from alkalines.
- Then there is leakage, and that is the one that fills your inbox with warranty claims. Alkaline cells leak potassium hydroxide when they sit in a device too long, and that corrodes your contacts and destroys the product from the inside. Your customer does not blame the battery brand for that, they blame you.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
Low self-discharge NiMH cells	A solid swap for high-drain products. Remember they are 1.2 V per cell instead of 1.5 V, so design your voltage range around that.
Lithium iron disulfide cells (L91 class)	For the toughest conditions, these handle heavy current and extreme cold far better than any alkaline will, and they do not leak.
A rechargeable lithium cell with protection	If the product is charged regularly anyway, a single protected lithium cell removes the chemistry problem entirely.

BEFORE YOU COMMIT

- **Peak current, measured:** Measure the product's peak current, including radio transmit bursts and motor starts, and compare it with the cell's pulse discharge curve, not its capacity rating.
- **Cutoff voltage:** Set the brown-out threshold and the low-battery warning around the loaded voltage of the cell you chose, which for NiMH is about 1.0 V per cell.
- **Leak protection:** If alkalines are still a possibility, add a low-battery warning that tells the customer to remove dead cells, and use gold-plated contacts.

Design for the chemistry you chose

NiMH, lithium primary, and alkaline all have different voltage curves. Whichever one you pick, the regulator, the brown-out threshold, and the fuel gauge all have to be designed around that curve.

5.3

Coin Cells for Anything Beyond Tiny Loads

PRIMARY CELL

WHY IT FAILS IN A PRODUCT

- Coin cells look like free energy when you are laying out a small product, because a CR2032 is tiny, cheap, gives you 3 V, and the datasheet promises around 225 mAh. That number only holds if you sip current out of it very slowly, a few hundred microamps or less.
- A coin cell has internal resistance in the tens of ohms when it is fresh, and that climbs into the hundreds of ohms as the cell drains. So the moment your Bluetooth radio fires off a transmit burst and pulls 10 or 15 mA, the voltage collapses and your microcontroller resets.
- You end up with maybe a third of the rated capacity, and the product dies early in the field for reasons that never showed up on your bench with a fresh cell.
- There is a legal side now too. Reese's Law passed in 2022 in the US, and any consumer product using a button or coin cell has to keep that battery in a compartment children cannot get into, plus carry specific warning labels on the product, packaging, and manual.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
A AA-sized lithium iron disulfide cell	If the product needs more than a trickle of current, a 1.5 V lithium AA in the L91 class gives you the same shelf life with far lower internal resistance.
A small rechargeable lithium cell	For products charged regularly, a small protected pouch cell handles radio bursts without a second thought.
A bulk capacitor across the coin cell	If a coin cell truly is the right choice, a large capacitor rides out the current spikes, but treat it as a patch on a battery that was never sized for the job.

BEFORE YOU COMMIT

- **Pulse current against the curve:** Compare the product's peak pulse current with the coin cell manufacturer's pulse discharge data, not the capacity number on the front page.
- **Reese's Law checklist:** If a coin cell stays, the compartment needs a tool or two independent and simultaneous hand movements to open, it has to stay closed through the rule's use and abuse tests, and the warning labels go on the product, the packaging, and the manual.
- **Real capacity test:** Run one unit on a fresh cell in the real duty cycle until it dies, and compare the result with the datasheet capacity before you promise a battery life.

An enclosure decision, not just a battery choice

Reese's Law turns a coin cell into an enclosure design requirement and a certification item, so the decision has to be made before the mechanical design is final.

5.4 Sealed-In Batteries

REGULATION

WHY IT FAILS IN A PRODUCT

- Sealing the battery in is one of the big mistakes I made with my own product, a small LED light on coin cells. A battery door meant another expensive mold, a bigger enclosure, and contacts that had to survive people jamming cells in backwards, on a product that sold for a few dollars, so I sealed it.
- Most customers still wanted to change the battery, and the complaints kept coming until I recorded a video showing them how to fully disassemble the unit to do it.
- Today that same decision would lock the product out of the EU. Article 11 of the EU Batteries Regulation says portable batteries have to be readily removable and replaceable by the end user, and it applies from February 18, 2027. Replacement has to work with commercially available tools, so adhesives that need heat or solvents are out, and you have to stock replacement batteries for five years after the last unit ships.
- This hits rechargeable products hardest, because a portable battery under that rule is anything sealed weighing 5 kg or less, which covers a glued-in lithium polymer pouch cell the same as a coin cell.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Decide it at the concept stage	This is an enclosure, tooling, and layout decision made at the beginning, and brutally expensive to reverse once you have cut steel.
Use a door held shut with a screw	A screwdriver counts as a commercially available tool and reusable fasteners are allowed, and a screw also satisfies the coin cell rule in entry 5.3.
Use a connector, not solder or glue	A pouch cell on a connector, held by a bracket rather than an adhesive that needs heat or solvents, is replaceable by the customer and by your repair line.
Confirm any exemption with an expert	Exemptions exist, and the latest round added wearables and electric toys, but confirm yours with a compliance expert before you commit to tooling.

BEFORE YOU COMMIT

- **Replacement stock plan:** Confirm you can buy and stock the exact battery for five years after the last unit ships, which turns a one-off pack design into a five-year commitment.
- **Tools test:** Have someone outside the team replace the battery using only a common screwdriver, and time it.
- **Two rules, one door:** If the product uses a coin cell, confirm the same screw-secured door satisfies both the EU rule and Reese's Law.

Sealed in, locked out

The rule is not about a component you can swap later. It is about the enclosure, and once the mold is cut, the cheapest fix becomes the most expensive one.

5.5 Unprotected 18650 Cells

LITHIUM-ION

WHY IT FAILS IN A PRODUCT

- The 18650 is an excellent lithium-ion cell, and that is exactly why so many people drop a bare one straight into their product. An unprotected cell has no circuitry between the chemistry and your customer, so nothing stops an over-discharge, an overcharge, or a dead short across the terminals.
- Take a lithium-ion cell below about 2.5 V and you cause permanent internal damage that can create a short weeks later, long after the unit shipped. A shorted 18650 can dump 10 A or more into whatever happens to be touching it, and that is how a product catches fire in somebody's living room.
- Your customers will find every one of those failure modes, since they will install cells backwards, charge them with whatever adapter is in the drawer, and toss a spare cell loose in a bag with their keys.
- There is a counterfeit problem on top of that, because the online marketplaces are full of rewrapped 18650s advertising 9,900 mAh when the best real cells top out around 3,500 mAh, as covered in entry 13.6.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
Protected cells	A protection circuit on the cell itself gives you overvoltage, undervoltage, overcurrent, and short-circuit protection right at the source, and it keeps you safer during development.
A proper pack with a battery management system	For multi-cell products, a pack with a BMS handles charge and discharge protection and cell balancing, with documentation you can hand to a test lab. The charger is still a separate part.
Cells from Samsung, LG, Molicel, or EVE	Build on cells from the top-tier makers, bought through authorized distribution, and get the real datasheet before you design anything around them.

BEFORE YOU COMMIT

- **Datasheet in hand:** You have the real cell datasheet, with the discharge curves, the temperature limits, and the cycle life, from the manufacturer rather than a seller.
- **Reverse insertion:** If the cell is user-replaceable, the holder or the circuit survives a cell installed backwards without damage, and entry 9.1 covers how.
- **Charger and cutoffs:** The charger IC and the protection thresholds match the specific cell's limits for voltage, current, and charging temperature.

Damage now, fire later

Lithium damage is cumulative and delayed. An over-discharge that happens today can become the internal short that starts a fire long after the unit shipped.

5.6 Bare LiPo Pouch Cells

LITHIUM POLYMER

WHY IT FAILS IN A PRODUCT

- Bare lithium polymer pouch cells are the most dangerous habit that carries over from the hobby bench into a product you sell. These are the flat silver packs with two wires hanging off them, the kind sold for RC planes and drones, and they ship with no protection circuit at all.
- That is on purpose, since RC users want every last amp out of the cell and accept that risk themselves. Your customer never agreed to any of that, and your name is the one on the box.
- The pouch is soft aluminum laminate, so a screw boss pressing into it, a pinched corner during assembly, or a drop onto a hard floor can puncture the cell.
- A punctured or overcharged lithium polymer cell does not just stop working. It swells, vents, and can go into thermal runaway.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
A pouch cell with an integrated protection board	From a supplier who will hand you a datasheet and test reports without you chasing them. On a protected cell that board usually sits at the top of the pouch, under the semi-transparent yellow tape where the wires come out.
A proper charge IC with temperature monitoring	Constant current and constant voltage charging with a hard limit at the cell's rated charge voltage, usually 4.2 V, and a thermistor input so charging stops when the cell gets too hot or too cold.
Mechanical design built around the cell	A little gap for swelling as the cell ages, and nothing sharp or rigid pressing against the pouch.

BEFORE YOU COMMIT

- **Protection board present:** Confirm the protection board is on the cell you will actually ship, not just on the sample, and that its cutoff values are in the datasheet.
- **Swelling allowance:** Leave clearance above the pouch for the thickness growth the cell datasheet specifies over its life.
- **Documentation:** Get the UN38.3 test summary and the cell datasheet before you send money, as entry 5.7 explains.

Look under the yellow tape

The quickest check on a pouch cell is the top edge. A protected cell has a small circuit board under the tape where the wires come out, and a bare RC cell has nothing there but the tabs.

5.7 Uncertified Lithium Packs

CERTIFICATION

WHY IT FAILS IN A PRODUCT

- This is the one that stops your entire product from shipping. These are the cheap packs sold online with no datasheet, no test reports, and a seller who cannot tell you which cells are inside.
- UN38.3 is the transport safety standard for lithium batteries, a set of eight tests covering altitude, thermal cycling, vibration, shock, external short circuit, crush, overcharge, and forced discharge. Without that test summary in your hands, your freight forwarder will not accept the shipment, and you cannot legally move those batteries by air.
- It blocks you at certification too, since IEC 62133 is the battery safety standard your product certification leans on, and you cannot get there with a cell you cannot trace back to an actual manufacturer.
- The liability side is even worse, because if one of your products burns down a customer's house and you cannot produce a single document about the battery inside it, you are the one holding all of it.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Buy from established pack manufacturers	Choose suppliers who hand you the full documentation package as a matter of course, because that is what a real battery vendor does.
Ask for the documents before you pay	Request the UN38.3 test summary, the cell datasheet, and the IEC 62133 report before you send anyone money, and treat a vague answer as a no.
Match the cells to the documents	Confirm the cell model in the test summary is the cell model in the pack you receive, since a substituted cell voids every report.

BEFORE YOU COMMIT

- **UN38.3 test summary:** In hand, naming the exact cell and pack model, with the test lab and the dates.
- **IEC 62133 report:** In hand for the cell or the pack, because your product safety certification will ask for it.
- **Cell traceability:** The cell manufacturer and model are named in writing, and they match a top-tier maker you can look up.

The big red flag

If a supplier gets vague the moment you ask for those documents, that is your answer. A vendor who has the paperwork sends it in minutes.

5.8 Custom Battery Packs and Chemistries

TECHNOLOGY TO AVOID

WHY IT FAILS IN A PRODUCT

- One of the most common reasons people choose a custom battery is that they are trying to squeeze in every last bit of capacity. If you are building something small, like a wearable or a handheld, it is tempting to design a battery that perfectly fits the internal shape.
- The moment you design a custom pack, you are responsible for meeting strict safety and transportation standards yourself. The testing is expensive, time-consuming, and failure-prone, especially for small teams without dedicated compliance engineers.
- You will face delays if your certification lab finds anything questionable, and you will have to source specialized suppliers who can build packs at low volumes while still supporting traceability and quality control.
- An unusual chemistry adds a second layer of the same problem, because your charger design, your protection thresholds, and your certification all have to be built for it from scratch.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Use a pre-certified off-the-shelf pack	A standard lithium polymer pack from a real manufacturer comes with the UN38.3 summary and the IEC 62133 report already done.
Design the enclosure around the pack	Pick the battery early and shape the enclosure to it, rather than shaping the battery to an enclosure that was designed first.
Accept a little unused volume	You will not use 100 percent of the internal volume, but you will avoid a year-long certification slog and a major liability risk if something goes wrong.

BEFORE YOU COMMIT

- **Standard sizes first:** Search the catalogs of established pack makers for the nearest standard size before anyone draws a custom outline.
- **Certification cost and time:** Get a quote and a schedule for UN38.3 and IEC 62133 on a custom pack, and put both next to the capacity you would gain.
- **Low-volume supplier:** If custom is unavoidable, confirm a pack builder will make your quantities with traceable cells and full documentation before you design around it.

Capacity is cheaper than certification

The few percent of volume a custom pack recovers is almost never worth the certification, the liability, and the supplier risk that come with it.

5.9

TP4056 Charging and Protection Boards

CHARGING

WHY IT FAILS IN A PRODUCT

- Those little charging modules built around the TP4056 are everywhere, they cost about fifty cents, and they are handy to have on a bench. The trouble is what is actually on them.
- The charger chip is often a clone, the charge current gets set by a single resistor the seller picked with no idea what battery you are using, and the protection circuit is sometimes missing entirely, with bare pads where the protection chip and the two switches that cut off the battery belong.
- That means you can end up pushing 1 A into a 400 mAh cell because the board came configured for a much larger battery, with nothing watching for an over-discharge or a short.
- Most of these boards also have no thermistor input, so charging continues at any temperature, which is exactly the condition that damages a lithium cell.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
Microchip MCP73831 or MCP73833	Cheap, well-documented single-cell chargers where you set the charge current with one resistor you chose. The 73831 has no thermistor input, so use the 73833 or the BQ series when you need temperature monitoring.
TI BQ series (such as the BQ24072 or BQ25601)	My favorite family, with power path management, thermistor monitoring, and options for USB input current limiting.
A dedicated protection IC or a protected cell	Protection belongs at the cell or on your board, as a separate function you can verify, never as an optional footprint on a fifty-cent module.

BEFORE YOU COMMIT

- **Charge current matches the cell:** Set the charge current to the cell datasheet's recommended rate, usually 0.5C or less, and confirm it with a meter on the first build.
- **Thermistor connected:** The charger's temperature input is wired to a thermistor on the cell, and charging stops outside the cell's charging temperature window.
- **Termination verified:** Confirm on the bench that charging terminates at 4.2 V and the taper current, rather than trickling forever.

Bench tool, not a product part

A TP4056 module is a fine way to charge a cell on your desk. In a product it is a lithium cell running with no safety net, configured by a stranger.

5.10 Wireless Charging (Qi and Others)

TECHNOLOGY TO AVOID

WHY IT FAILS IN A PRODUCT

- Wireless charging feels futuristic, and for marketing it sounds great. Just drop the device on a pad and it charges. Under the hood it is a different story.
- Coil alignment has to be right, thermal issues come up fast, and efficiency is lower than wired charging, so the product gets warm while it charges slower. If the power transfer is inconsistent, customers start complaining.
- Qi certification is not cheap, and inductive charging introduces a whole new set of EMI issues that are hard to troubleshoot, on top of the wireless certification the product already needed.
- The receiver coil also dictates the enclosure, since it needs a non-metal back, a defined spacing to the pad, and room away from the battery and the antenna.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
A USB-C port	The default for most consumer products, with chargers customers already own and the reasons covered in entry 8.1.
Pogo pins	For sealed or waterproof products, spring-loaded contacts give you a clean charging interface with no open port.
A custom cradle	A dock with contacts still feels sleek, and it is far easier to implement and much more reliable than a coil.

BEFORE YOU COMMIT

- **Is it a must-have:** Ask whether a single customer would choose a competitor because your product has a port instead of a pad. If not, it is not a must-have.
- **Thermal budget:** If wireless stays, model the heat from a receiver running at typical efficiency inside the sealed enclosure, next to the battery.
- **Certification budget:** Price Qi certification and the extra EMI work, and add them to the schedule before the feature goes on the spec sheet.

Only when it is the point of the product

Wireless charging earns its cost when it is the reason the product exists, like a fully sealed device that cannot have contacts. Otherwise a port or a cradle wins on cost, heat, and reliability.

STAGE: POWER DESIGN

VOLTAGE REGULATORS AND POWER

These regulators show up all the time in beginner designs and early prototypes. They work fine on the bench, but in a real product they cause thermal failures, blown EMC tests, and production problems that cost thousands of dollars to fix. I spent years at Texas Instruments designing voltage regulators, including several popular ones, and every part in this section has a better, modern replacement that costs about the same.

IN THIS PART

6.1	The 7800 Series (7805 and Friends)	p. 53
6.2	LM2596	p. 54
6.3	MC34063	p. 55
6.4	XL6009, XL4015, and Other No-Name Switchers	p. 56
6.5	AMS1117	p. 57
6.6	Cheap, Unshielded Switching Regulator Modules	p. 58
6.7	HT7333 and HT7533	p. 59
6.8	LM317 and Adjustable Regulators Set to a Fixed Voltage	p. 60
6.9	MCP1700 Chosen on Quiescent Current Alone	p. 61
6.10	Unshielded DC-DC Converter Inductors	p. 62

6.1 The 7800 Series (7805 and Friends)

LINEAR REGULATOR

WHY IT FAILS IN A PRODUCT

- The 7800 series is the first regulator family every engineer learns about, and it feels like a safe, proven choice because it has been around for over 50 years.
- It has a dropout voltage of around 2 V, which means the input must always be at least 2 V higher than the output, and with any linear regulator all the power across that voltage difference gets wasted as heat.
- Say your input is only about a volt above your desired output, like running a 5 V rail from a 6 V battery pack. A modern LDO handles that easily. The 7800 series cannot even regulate in that situation, so you would need at least 7 V on the input just to get a stable 5 V out, and every extra volt is heat.
- It also comes in large, outdated packages that eat up board space, and its quiescent current of several milliamps is a constant drain that a battery product cannot afford.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
A synchronous buck converter	For any big voltage step-down, a buck converter steps voltage down far more efficiently than any linear regulator can, and entry 6.2 covers what to look for.
A modern LDO with real low dropout	For small step-downs where a linear regulator makes sense, choose a part with dropout under a few hundred millivolts at your load current.
A fixed-output LDO matched to the rail	Pick the exact output voltage you need, in a modern SOT-23 or DFN package, and keep the adjustable parts for the reasons in entry 6.8.

BEFORE YOU COMMIT

- **Power dissipation, calculated:** Multiply the input-to-output voltage difference by the load current. If the result is more than a few hundred milliwatts in a small package, a linear regulator is probably the wrong topology.
- **Dropout at your current:** Read the dropout spec at your actual load current and temperature, not the headline number, because dropout climbs with both.
- **Quiescent current:** Check the quiescent current against the product's sleep budget, since the regulator draws it whether the load is asleep or not.

Leave it in the classroom

The 7800 series is a fine teaching part. In a product, every volt of headroom it demands is heat you have to get rid of and battery life you gave away.

6.2

LM2596

SWITCHING REGULATOR

WHY IT FAILS IN A PRODUCT

- The LM2596 is everywhere because of those one-dollar buck converter modules. Plug it in, get your voltage, and move on, which makes it great for quick prototyping and terrible for a product.
- It switches at just 150 kHz, which means it needs large inductors and large capacitors, and it throws off a ton of electromagnetic interference. The footprint is huge compared to modern alternatives.
- The module layouts were never designed for EMC compliance, so you will fail FCC and CE testing, and the unshielded inductor on most of them makes it worse, as entry 6.10 explains.
- Counterfeit LM2596 chips are all over the market, so you cannot even trust the datasheet specs on the part you actually receive.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
A modern synchronous buck converter	Look for a part switching at 500 kHz or higher from TI, Analog Devices, ST, or Monolithic Power Systems, with integrated switches, smaller passives, and far less EMI.
Your own layout from the reference design	Always design your own layout using the IC vendor's reference design rather than dropping in a random module, and follow entry 10.2 for the hot loop.
A quality shielded module if you must use one	If you want a module, use one from a reputable supplier that is shielded and properly laid out, as covered in entry 6.6.

BEFORE YOU COMMIT

- **Switching frequency:** Prefer a part switching at 500 kHz or above for smaller passives, and check that the frequency and its harmonics do not land on a band your radio or your sensors care about.
- **Synchronous rectification:** Prefer parts with an integrated synchronous switch rather than an external Schottky diode, for efficiency and for a smaller hot loop.
- **Datasheet source:** Buy the regulator from an authorized distributor so the datasheet you designed to is the part you receive.

Cost check

A modern buck converter IC and its passives cost about what the module does, and you get a smaller footprint, a quieter board, and a chance of passing EMC on the first try.

6.3

MC34063

SWITCHING REGULATOR

WHY IT FAILS IN A PRODUCT

- The MC34063 can technically do buck, boost, and inverting topologies all in one chip, which makes it look like the ultimate flexible power solution on paper. In practice it is a headache.
- It uses a gated-oscillator control scheme, so the switching frequency wanders with the load and the EMI is broadband and hard to filter, and efficiency in a real design usually lands around 80 percent or lower.
- The external component count is high, the output ripple is large, and most designers end up with a noisy output without knowing why.
- Debugging a bad MC34063 design can eat up weeks of engineering time, and I have seen that happen more than once.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
A dedicated synchronous buck IC	For step-down, a purpose-built converter outperforms the MC34063 in every way with a simpler, cleaner design and a proper reference layout.
A dedicated boost IC	For step-up, modern boost converters from TI, Analog Devices, or ST come with internal compensation and integrated switches, so the design is a handful of passives.
An inverting or buck-boost IC where needed	If the product needs a negative rail or an input that crosses the output voltage, a purpose-built buck-boost part handles it without the guesswork.

BEFORE YOU COMMIT

- **One topology per part:** Write down what the rail actually needs to do, then choose a part designed for exactly that, rather than one that can do everything badly.
- **Efficiency at your load:** Read the efficiency curve at your real load current and input voltage, since the peak number in the summary rarely matches your operating point.
- **Reference design available:** Confirm the vendor publishes a reference layout, and copy it.

Flexible is not the same as good

A chip that can do three jobs usually does none of them well. Purpose-built converters cost the same and take a fraction of the engineering time.

6.4

XL6009, XL4015, and Other No-Name Switchers

SWITCHING REGULATOR

WHY IT FAILS IN A PRODUCT

- No-name switcher ICs like the XL6009 and XL4015 show up on ultra-cheap modules, and they seem to work fine during testing.
- The datasheets for these parts are often incomplete, inconsistent, or outright fabricated. There is no way to verify the electrical specs on your own, there is no factory application support, and if you hit a design issue there is nobody to call.
- Long-term sourcing is also unreliable since these parts appear and disappear from the market without any warning, so the regulator in your second production run may not exist.
- Without a trustworthy datasheet, you cannot do the thermal calculation, the stability analysis, or the EMI work that a switching regulator needs, so every one of those becomes a guess.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
Regulators from TI, Analog Devices, or ST	You get verified datasheets, reference designs, application engineers you can actually talk to, and reliable distribution.
Monolithic Power Systems or Richtek for cost-sensitive rails	Both publish real datasheets and are stocked by major distributors, which is the test that matters.
A rule for the whole bill of materials	If you cannot buy it from a major distributor like Mouser or Digi-Key, it does not belong in your product.

BEFORE YOU COMMIT

- **Distributor stock:** Search the exact part number on two major distributors. If neither stocks it, stop.
- **Datasheet quality:** Look for a full electrical characteristics table with minimum and maximum values, typical performance curves, and a layout guideline section. A two-page summary is not a datasheet.
- **Application support:** Confirm the vendor has an application forum or support channel where a real engineer answers questions.

Cost check

The per-unit cost difference between a no-name switcher and a part from a major vendor is often under fifty cents, and that is the cheapest application support you will ever buy.

6.5

AMS1117

LINEAR REGULATOR

WHY IT FAILS IN A PRODUCT

- The AMS1117 is on nearly every Arduino, ESP32, and STM32 dev board out there, so designers assume that if it works on the dev board it will work just fine in their product. That is a dangerous assumption.
- Despite being marketed as low dropout, the actual dropout voltage is over 1 V at rated current, which is not low at all, and it wastes that headroom as heat.
- Transient response is poor, so current spikes from Wi-Fi transmissions or motor startups cause voltage sags that can reset your microcontroller.
- Counterfeit AMS1117 chips are everywhere, and the fakes perform even worse than the originals, so the behavior changes from batch to batch.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
A modern LDO with dropout under 250 mV	Real low dropout lets you run a smaller input-to-output gap, and that gap is where the heat and the battery life go. At the same input voltage the heat is the same.
TI TPS793 (200 mA)	One of the first chips I designed at TI, with high power supply rejection and low noise for sensitive analog and RF rails up to 200 mA.
TI TPS799 (200 mA)	The same kind of performance with lower quiescent current, for low-current rails in products that spend most of their life asleep. A 1 A AMS1117 rail needs a bigger part.

BEFORE YOU COMMIT

- **Load transient in the datasheet:** Look at the load transient response plot and compare the dip with what your microcontroller's brown-out threshold will tolerate.
- **Dropout at your current and temperature:** Read the dropout curve at your peak load, since a part that is fine at 100 mA may be out of regulation at 500 mA on a hot day.
- **Quiescent current in sleep:** Confirm the regulator's own draw fits the product's sleep budget.

Dev boards are not reference designs

A dev board is built to be cheap and to work on a desk with a wall adapter. The regulator on it was chosen for those reasons, not for your battery life or your EMC test.

6.6

Cheap, Unshielded Switching Regulator Modules

POWER MODULES

WHY IT FAILS IN A PRODUCT

- There is nothing wrong with using a well-designed switching regulator module in a product, and I actually encourage it in a lot of cases because it saves you from having to do your own switcher layout. The problem is grabbing the cheapest module you can find and dropping it straight into a production design.
- Those cheap modules are almost never shielded, which means they radiate EMI everywhere and you will have a very hard time passing FCC or CE testing.
- The PCB layout on those modules is usually terrible too, which kills the switching IC's performance and makes thermal management unpredictable.
- Quality varies from batch to batch, so the power supply that worked in your first 100 units might behave differently in the next 500.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
A shielded module from a reputable supplier	Look for a module that is properly shielded, uses a quality layout, and has consistent sourcing, from vendors like Murata, TDK, CUI, Recom, or Traco.
A module with a real datasheet and EMC data	A serious module vendor publishes efficiency curves, thermal derating, and often conducted and radiated emissions plots.
Your own layout from the IC vendor's reference design	If your team has the layout experience, an on-board regulator from a major vendor is usually cheaper at volume.

BEFORE YOU COMMIT

- **Shield present:** The module has a metal can or a shielded inductor, and the vendor says so in the datasheet.
- **Derating curve:** Check the output current derating against the temperature inside your enclosure, not the room.
- **Batch consistency:** Confirm the vendor controls the bill of materials on the module, so the second batch matches the first.

When a module is the smart choice

A quality module can be exactly the right call, especially if you do not have switching regulator layout experience on your team. The mistake is the cheapest one, not the category.

6.7

HT7333 and HT7533

LINEAR REGULATOR

WHY IT FAILS IN A PRODUCT

- The HT7333 and HT7533 cost just a few pennies per unit, which makes them look like a no-brainer for simple low-power circuits.
- The maximum rated current is only 250 mA on the HT7333 and just 100 mA on the HT7533, and both ratings drop further at elevated temperatures.
- There is no thermal shutdown protection, so either part can overheat and damage itself without giving you any warning.
- Load regulation falls apart above 100 mA, and power supply rejection drops off above 1 kHz, which means switching noise from other regulators on your board passes right through to your sensitive circuits. That kind of noise coupling causes strange behavior in a microcontroller or a sensor, and it is extremely hard to track down if you do not know the regulator is the source.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
An LDO with thermal shutdown and current limit	For a few cents more per unit you get a part that protects itself and the load when something goes wrong.
An LDO with power supply rejection specified to 100 kHz or beyond	That is the range where switching noise lives, and it is where the HT parts stop rejecting anything.
A part specified across temperature	Choose a regulator whose datasheet gives load regulation and current limits at the full temperature range, not just at 25 degrees C.

BEFORE YOU COMMIT

- **Thermal shutdown:** The datasheet lists a thermal shutdown temperature. If it does not, the part has none.
- **Power supply rejection plot:** The PSRR curve is shown out to at least 100 kHz, and it is still useful at your switching regulator's frequency.
- **Current headroom:** Your peak load, including radio bursts, sits below half the rated current at the hottest temperature the enclosure reaches.

Cost check

For just a few cents more per unit you can get an LDO with thermal protection, better load regulation, and specs you can actually rely on across temperature.

6.8

LM317 and Adjustable Regulators Set to a Fixed Voltage

LINEAR REGULATOR

WHY IT FAILS IN A PRODUCT

- The LM317 is the most common adjustable regulator, and it feels like the ultimate flexible part because one chip can give you any voltage. The problems here apply to any adjustable regulator.
- The external resistor divider adds tolerance stackup, where small errors in each resistor compound, plus thermal drift and two extra failure points. On the LM317, if the bottom resistor opens the output jumps to the full unregulated input voltage and fries everything downstream, and if the top resistor opens the output drops to about 1.25 V. Other adjustable parts fail differently, but the divider is still two more ways to fail.
- The LM317 specifically has no enable pin and a dropout over 2 V. It also has no ground pin, so it is a floating topology where the chip's own supply current gets dumped straight into the output, which is why it needs a minimum load of 3.5 to 10 mA just to stay in regulation. Draw less and the output rises.
- Almost everyone using an adjustable regulator is setting it to a fixed voltage anyway.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
A fixed-output LDO matched to your rail	Fewer external components, better specs, lower risk, and no resistor divider to fail open.
A fixed LDO with an enable pin	Lets the microcontroller shut the rail down in sleep, which the LM317 cannot do at all.
An adjustable part only for a truly variable rail	If the product really needs a programmable voltage, use a modern adjustable LDO with a ground pin, and use 0.1 percent resistors in the divider.

BEFORE YOU COMMIT

- **Minimum load:** If any adjustable regulator stays, confirm the load never drops below the minimum load current, including in sleep.
- **Divider failure modes:** Ask what happens downstream if either divider resistor opens, and confirm the answer is not a dead product.
- **Enable and sleep:** Confirm every rail that should shut down in sleep has a regulator with an enable pin.

The nasty surprise in low-power designs

A regulator with a minimum load requirement is a regulator that misbehaves exactly when your product is trying to save power, and that is the state it spends most of its life in.

6.9

MCP1700 Chosen on Quiescent Current Alone

LINEAR REGULATOR

WHY IT FAILS IN A PRODUCT

- The MCP1700 has a quiescent current of just 1.6 uA, and for battery-powered products that number is almost irresistible.
- The maximum output current is only 250 mA, and the transient response is very slow. When your microcontroller wakes from deep sleep and draws a burst of current, the MCP1700 cannot respond fast enough and the output voltage dips hard.
- If your microcontroller needs 150 mA or more in short bursts, you will see brown-out resets that are incredibly frustrating to debug, because everything looks perfectly fine under steady-state testing.
- People pick this part based entirely on the quiescent current spec without ever checking whether it can handle their real load profile, and a single spec on a datasheet never makes a regulator right for a product.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
An LDO matched to the full load profile	Look for low quiescent current, fast transient response, and enough current capacity for your worst-case bursts, all three at once.
TI TPS7A02 (200 mA) or similar nanopower LDOs	Modern parts get quiescent current into the tens of nanoamps with fast transient response, so you no longer trade one for the other, but check the current rating against your bursts.
A bulk capacitor sized for the burst	Whatever LDO you choose, a properly sized output capacitor near the load covers the first microseconds of a transmit burst.

BEFORE YOU COMMIT

- **Wake-up burst, measured:** Capture the current draw on a scope as the microcontroller wakes and the radio transmits, and use that peak, not the average, to size the regulator.
- **Transient response plot:** Compare the datasheet's load step response with your brown-out threshold, and leave margin for a cold battery.
- **Sleep current, whole board:** Measure the sleep current of the whole board, because the regulator is rarely the biggest contributor once the design is done.

Look at the full picture

A single spec on a datasheet does not make a regulator right for your product. Match the regulator to your full load requirements, not to one headline number.

6.10 Unshielded DC-DC Converter Inductors

SWITCHING REGULATOR

WHY IT FAILS IN A PRODUCT

- Switching regulators are everywhere in product designs, but pairing one with an unshielded inductor is asking for EMI trouble.
- The magnetic field from an unshielded inductor radiates freely, couples into nearby traces, and contaminates sensitive analog or RF sections of your board.
- This often shows up as noise in a pre-compliance scan, or worse, as an outright failure at the test lab when you are paying thousands of dollars for certification testing.
- It also degrades the product itself, because the coupled noise ends up in your ADC readings, your audio, and your radio receiver sensitivity long before it fails a test.

SWAP IT FOR

REPLACEMENT PART	WHY THIS PART
A shielded inductor from Wurth, Coilcraft, or TDK	Shielded parts cost only slightly more per unit and reduce radiated EMI significantly.
A semi-shielded inductor for cost-sensitive rails	A middle ground that contains most of the field for a few cents less than a fully shielded part.
Good layout around the regulator	The inductor is half of the fix. Keeping the hot loop tiny, as entry 10.2 covers, is the other half.

BEFORE YOU COMMIT

- **Shielded in the part number:** Confirm the inductor series is shielded or semi-shielded in the datasheet, since the words rarely appear in the part number itself.
- **Saturation current:** Check the saturation current against the regulator's peak inductor current with margin, because a saturating inductor is noisy and hot.
- **Distance to sensitive circuits:** Keep the inductor away from the antenna, the analog front end, and any audio path, and never on the opposite side of the board directly beneath them.

Cost check

A tiny saving on the bill of materials can easily turn into a 5,000 dollar or more certification failure if your design does not pass EMC testing.

STAGE: SCHEMATIC DESIGN

CAPACITORS, PASSIVES, AND DISCRETES

Capacitors are the most common part in your product and the part nobody thinks about. They rarely get a second look, so the same mistakes show up in schematic after schematic that I review. One of them will fail safety certification outright, three of them make a product flaky in ways you can never reproduce on the bench, and one of them can catch fire when it fails. The last two entries cover the other passive and discrete habits that follow people from the hobby bench into production.

IN THIS PART

7.1	Y5V and Z5U Ceramic Capacitors	p. 64
7.2	Ultra-Small Packages When You Are Not Space Constrained	p. 65
7.3	A Ceramic Capacitor on an LDO That Needs ESR	p. 66
7.4	X7R or X5R Where a COG Belongs	p. 67
7.5	General-Purpose Electrolytics Next to Anything Hot	p. 68
7.6	A Regular Capacitor Where an X or Y Capacitor Belongs	p. 69
7.7	Solid Tantalum Capacitors Without Heavy Derating	p. 70
7.8	Through-Hole Parts When SMD Alternatives Exist	p. 71
7.9	Mechanical Relays When Solid-State Works	p. 72

7.1 Y5V and Z5U Ceramic Capacitors

DIELECTRIC

WHY IT FAILS IN A PRODUCT

- If you have ever bought one of those cheap capacitor assortment kits, there is a good chance most of it was Y5V or Z5U. Those letter codes tell you how the capacitor behaves over temperature, and for these two the answer is badly.
- A Y5V capacitor is allowed to lose up to 82 percent of its capacitance across its rated temperature range, and Z5U is not much better at 56 percent. So that 10 uF cap on a regulator output might be acting like a 2 uF cap inside a hot enclosure.
- These dielectrics also lose a big chunk of capacitance from DC bias, covered in entry 7.2, so the two effects stack on top of each other.
- That means your bulk filtering and decoupling capacitors end up specified around a number the circuit never actually sees, and the regulator that was stable on the bench oscillates in the field.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
X7R for general decoupling and bulk capacitance	X7R stays within plus or minus 15 percent across its whole temperature range, which you can actually design around.
X5R where the temperature range allows it	The same tolerance up to 85 degrees C, often in a smaller package for the same value, which is fine for most consumer products.
C0G for anything small and precise	For timing, filters, and analog paths, the class 1 dielectric covered in entry 7.4 is the only right answer.

BEFORE YOU COMMIT

- **Dielectric on every capacitor line:** Put the dielectric code on the bill of materials for every ceramic capacitor, so an assembler cannot substitute a Y5V for an X7R.
- **Temperature of the enclosure:** Know the peak temperature inside the product and confirm the capacitor's rated range covers it with margin.
- **Assortment kits stay on the bench:** Never let a part from an assortment kit make it onto a production bill of materials without a manufacturer part number behind it.

Cost check

X7R costs a little more than Y5V, usually a fraction of a cent per part, for a capacitor that holds its value across the whole temperature range instead of losing most of it.

7.2

Ultra-Small Packages When You Are Not Space Constrained

PACKAGE AND DC BIAS

WHY IT FAILS IN A PRODUCT

- Product creators love the smallest capacitor package they can find, because smaller feels more professional. But an 0402 is only 1 mm long and an 0201 is 0.6 by 0.3 mm, basically a speck of dust. Hand soldering an 0201 is nearly impossible, and reworking one on a finished board takes a microscope.
- Even in mass production, the smaller the part, the more you invite defects like tombstoning, where one end lifts off its pad during reflow, and solder bridging between pads that are almost touching. Some lower-cost assembly houses will not place 0201 parts at all, or charge extra for them.
- Squeezing a large capacitance into a tiny package also hurts the circuit electrically, through DC bias. A 10 uF ceramic in an 0402 was measured with almost no voltage across it. Put a normal 3.3 V rail on it and it can lose more than half of its value, so your 10 uF cap is doing the job of a 3 or 4 uF cap.
- The thinner the dielectric layers, the worse the effect, so the smallest package with the biggest value is the worst case, right where a switching regulator or a bulk filter counted on the full number.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Make 0603 the default	For most general-purpose capacitors, 0603 is a good default, and you only go smaller when you truly need the space or the better high-frequency performance.
Check the DC bias curve at your operating voltage	Murata, TDK, and Samsung all publish DC bias data and have free online tools that show the curve for most of their capacitors.
Move up a package or a part number	If you are losing more than you can live with, move to an 0603 or 0805, or try a different part number with a better curve.

BEFORE YOU COMMIT

- **Effective capacitance, not nominal:** For every bulk and regulator capacitor, write down the capacitance at the actual rail voltage from the vendor tool, and design to that number.
- **Assembler capability:** Ask your assembly house for their minimum package size and any surcharge before you choose 0201 anywhere.
- **Rework plan:** Confirm your own team can rework the smallest package you chose, because prototypes always need it.

Where DC bias bites hardest

A switching regulator output or a bulk filter that counted on a certain capacitance for ripple, transient response, or stability is where a capacitor that lost half its value shows up as a product that resets.

7.3

A Ceramic Capacitor on an LDO That Needs ESR

REGULATOR STABILITY

WHY IT FAILS IN A PRODUCT

- Every capacitor has a little resistance inside it called ESR, equivalent series resistance. Most of the time you want that as low as possible, but some older linear regulators actually depend on that resistance to stay stable. Their control loop was designed when the output capacitor was always a tantalum or an electrolytic, and the ESR in those parts acts like a built-in damper.
- Put a ceramic capacitor on one of those regulators, with almost no ESR, and the loop can start oscillating, which shows up as a noisy output, a regulator that runs hot, or a microcontroller that resets for no obvious reason.
- The classic example is the LM1117, which asks for an output capacitor with an ESR between 0.3 and 22 ohms, and a typical ceramic sits far below that.
- When I was designing LDOs at TI, almost all of them were built to be stable with a low-ESR ceramic output, and most newer, higher-quality LDOs are the same, but plenty of older designs are still out there.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Check the datasheet for the exact part	If the regulator needs ESR, give it a traditional tantalum or an aluminum electrolytic with an ESR inside the range the datasheet calls for.
Pick an LDO that is stable with ceramics	Any time the load current jumps, that current flows through the ESR and creates a dip in the output, so a regulator that depends on ESR gives up transient response to get its stability. If that dip matters, choose a ceramic-stable LDO instead.
Read the load transient plot	Before you commit to either regulator, check its load transient response in the datasheet against the current steps your product produces.

BEFORE YOU COMMIT

- **Output capacitor section read:** You have read the output capacitor requirements for the exact regulator part number, including the ESR range and the minimum value.
- **ESR at your frequency:** The ESR of the capacitor you chose sits inside that range at the temperature extremes, since electrolytic ESR climbs sharply in the cold.
- **Scope check on the first build:** Look at the regulator output on a scope under a load step. Ringing means the loop is marginal.

The hidden trade

A regulator that needs ESR to be stable is trading transient response for stability. That is fine for a slow load and a problem for a radio that wakes up in microseconds.

7.4 X7R or X5R Where a C0G Belongs

PRECISION ANALOG

WHY IT FAILS IN A PRODUCT

- X7R and X5R are class 2 ceramics, built to pack a lot of capacitance into a small size, and they give up accuracy to get there. For decoupling and bulk filtering, that trade is exactly what you want.
- But there is a whole category of circuits where they will wreck your design, and that is anything where the exact capacitance matters: RC timing, oscillators, crystal load capacitors, active filters, and anything in a precision analog signal path.
- X7R does not hold a value, it holds a range. You start with a tolerance of 10 or 20 percent, add about 15 percent of drift over temperature, then the DC bias loss, then aging, because these dielectrics slowly lose capacitance just sitting on the shelf. Stack all of that up and a filter corner or a timing interval can move by 30 percent or more.
- Class 2 ceramics are also piezoelectric, so vibration or even a tap on the board shows up as a voltage spike in your signal.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
C0G, also called NP0	A class 1 ceramic built for accuracy instead of size. It is stable to within about 30 ppm per degree C, it has no DC bias effect, no piezoelectric effect, and it barely ages.
C0G in larger packages for larger values	You can get it up to around 100 nF in small packages and higher in larger ones, which covers almost every timing and filter circuit.
X7R only where you need bulk	Use X7R where you need bulk capacitance, and C0G anywhere the exact value matters, and mark the difference on the schematic.

BEFORE YOU COMMIT

- **Circuit audit:** Walk the schematic and list every capacitor that sets a frequency, a time constant, a gain, or a crystal load. Every one of them is C0G.
- **Tolerance stack:** For any remaining class 2 part in a sensitive spot, add up tolerance, temperature drift, DC bias, and aging, and confirm the circuit still works at both ends.
- **Microphonics:** If the product has a sensitive analog path near a speaker, a motor, or a button, test it for vibration-induced noise before you blame the amplifier.

Cost check

In the picofarad to low nanofarad range where most timing and filter circuits live, C0G only costs a few cents more than X7R.

7.5

General-Purpose Electrolytics Next to Anything Hot

LIFETIME

WHY IT FAILS IN A PRODUCT

- Aluminum electrolytic capacitors have a lifespan, and it is shorter than most people think. The cheap general-purpose ones might have only a 2,000-hour endurance rating at 85 degrees C.
- An electrolytic has liquid inside it, and that liquid slowly evaporates through the seal, faster when it is hot. Every 10 degrees C increase in temperature roughly cuts the life in half. At 85 degrees C, that 2,000-hour rating is less than three months of continuous operation. At 45 degrees C the same rule gives about 32,000 hours, a little under four years of always-on use.
- Now think about where these capacitors end up on a lot of boards: right next to a hot linear regulator, a power resistor, or a switching transistor, inside an enclosure with no airflow.
- When one finally dries out, the ESR climbs, the ripple gets worse, and you get a product that works for a year and then starts rebooting or humming.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
105 C parts with a 5,000-hour rating or better	This usually adds a few cents at most, and it multiplies the life at any given temperature several times over.
Layout that keeps them away from heat	Electrolytics are one of the few parts that wear out, so plan for that from the start and keep them away from regulators, power resistors, and switching transistors.
Ceramic or solid polymer aluminum	In a lot of low-voltage circuits that do not depend on ESR, replace the electrolytic with a ceramic, or with a solid polymer aluminum, which still ages but has no liquid to dry out.

BEFORE YOU COMMIT

- **Life calculation:** Estimate the capacitor's core temperature in the enclosure, then apply the rule of thumb to the datasheet rating and compare the result with the product's expected life.
- **Ripple current rating:** Confirm the ripple current in the circuit sits well below the capacitor's rating, since ripple heats the part from the inside.
- **Brand and source:** Buy from a brand like Panasonic, Nichicon, or Rubycon through an authorized channel, for the reasons in entry 13.4.

Do the math once

Every 10 degrees C roughly halves the life. A capacitor sitting at 75 degrees C next to a regulator lives one quarter as long as the same part at 55 degrees C an inch away.

7.6

A Regular Capacitor Where an X or Y Capacitor Belongs

MAINS SAFETY

WHY IT FAILS IN A PRODUCT

- If your product plugs into the wall, there are two spots where a regular capacitor becomes a shock and fire hazard. The first is across the incoming AC line, between line and neutral, where a capacitor keeps noise from getting into or out of your product. The second is a capacitor that crosses from the AC side to the low-voltage side to reduce conducted emissions, right on the isolation barrier.
- A capacitor across the line sees every surge and spike the power grid throws at it, and a regular film or ceramic capacitor was never tested for that. If it fails as a short across the line, you get a fault current that can start a fire.
- If the one crossing the barrier fails as a short, mains voltage ends up on the low-voltage side of your product, which is the side your customer is touching.
- Put a standard capacitor in either of those spots and the product fails safety certification, no matter how well the rest of the design was done.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
An X-rated capacitor across the line	X capacitors are tested to survive those surges, and because they sit between line and neutral, the standard tests for the fire risk. For most consumer products that means an X2.
A Y-rated capacitor across the barrier	Y capacitors are built and tested so a short is extremely unlikely, since a short there puts mains on the outside of the product. Y1 is the grade rated to bridge the barrier on its own.
A pre-certified external adapter	If this sounds like a risk you would rather avoid, keep mains out of your enclosure entirely with a certified external supply, as entry 13.8 covers.

BEFORE YOU COMMIT

- **Safety approvals on file:** The manufacturer's safety approvals for the exact X and Y part numbers are in your compliance folder, because the certification lab will want to see them.
- **Subclass matches the application:** The X subclass matches the surge level for your installation, and the Y subclass matches the insulation grade of the barrier it crosses.
- **No substitutions:** Both parts are marked on the bill of materials as safety-critical with no substitutions allowed.

Cost check

An X2 across the line and a Y1 across the barrier together cost well under a dollar, and they are the difference between passing safety certification and failing it outright.

7.7

Solid Tantalum Capacitors Without Heavy Derating

FIRE RISK

WHY IT FAILS IN A PRODUCT

- This is the one capacitor with a fire risk baked into its chemistry, even in a low-voltage circuit with no mains anywhere near it. The traditional solid tantalum uses manganese dioxide for one of its internal layers, and that is an oxidizer.
- When one of these fails, it usually fails as a short circuit, and if enough fault current is available, that spot heats up until the oxidizer and the tantalum react and the part ignites. They often fail right at power-up, when the inrush surge hits a tiny weak spot in the dielectric.
- A 16 V tantalum on a 12 V rail looks fine on paper, but under a proper derating rule it is way over the line, and that is why some engineering teams have banned tantalums entirely.
- Polymer tantalums can still fail as a short, but without the oxidizer that short is far less likely to turn into a fire.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
Polymer tantalum capacitors	They fit the exact same footprint but replace the manganese dioxide with a conductive polymer, and depending on the part you can typically run them at 80 to 90 percent of rated voltage instead of 50 percent.
Heavy derating if you must use the traditional type	Never run them above half of the rated voltage, and on any rail that may see a surge, closer to a third, with some resistance in series to limit inrush.
A traditional tantalum only where it earns its place	Ultra-low-power standby circuits, where polymer leakage current matters, and the output of an old regulator that needs ESR, as in entry 7.3.

BEFORE YOU COMMIT

- **Voltage rating against the rail:** Every tantalum on the schematic is rated at least twice the rail voltage if it is manganese dioxide, or meets the vendor's derating guidance if it is polymer.
- **Inrush path:** There is resistance or a soft start between the supply and any traditional tantalum, so the power-on surge cannot hit it at full current.
- **Leakage in standby:** If the product sleeps on a battery, the polymer part's leakage current is in the sleep budget.

Cost check

Polymer tantalums cost more, usually somewhere in the tens of cents per part in volume, which is cheap insurance against a fire. If you truly must use a manganese dioxide part, treat it as the most dangerous component on your board, because it probably is.

7.8

Through-Hole Parts When SMD Alternatives Exist

ASSEMBLY

WHY IT FAILS IN A PRODUCT

- Through-hole parts are bigger, harder to place with automated assembly equipment, and they drive up manufacturing cost. People coming from the hobbyist world default to them because that is what they learned to solder by hand.
- Through-hole components require wave soldering or hand soldering as a separate step after SMD reflow, and that adds both cost and complexity to every single production run.
- Hand-soldered joints also carry more workmanship variation from unit to unit than the reflowed side of the board, unless you pay for inspection on them too.
- There are legitimate exceptions, like connectors that need mechanical anchor strength, some large power components, and certain high-reliability applications, but for the vast majority of resistors, capacitors, and ICs, SMD is the right choice.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Default to SMD for every passive and IC	If an SMD equivalent exists, and it almost always does, use it, so the board goes through the assembler's standard SMT process with no extra soldering step.
Keep through-hole for what needs it	Connectors that take mechanical stress, large electrolytics, transformers, and heavy power parts that need the anchor strength of a plated hole.
Ask your assembler about the second process	If any through-hole part remains, confirm whether the assembler will wave solder, selective solder, pin-in-paste reflow, or hand solder it, and what each costs per board.

BEFORE YOU COMMIT

- **Count the through-hole parts:** List every through-hole part on the board and justify each one with a mechanical or thermal reason. Habit is not a reason.
- **Single-side reflow:** Confirm the board can be built with SMD reflow on one or two sides and no separate soldering process, or know exactly what that extra process costs.
- **Prototype like production:** Build even the early prototypes with the SMD parts you will ship, so the layout and the footprints get tested along with the circuit.

One process, not two

Every extra soldering process is a line item on every unit you ever build. The cheapest board is the one that goes through the reflow oven once and comes out finished.

7.9

Mechanical Relays When Solid-State Works

SWITCHING LOADS

WHY IT FAILS IN A PRODUCT

- Mechanical relays wear out after a limited number of cycles, they are loud, they are physically large, and they draw significant coil current the entire time they are energized.
- For many products, especially battery-powered or compact designs, a solid-state alternative is the better choice, with no moving parts, no audible click, a fraction of the size, and dramatically longer life.
- Founders default to relays far too often simply because they are simple to understand, and then they pay for the coil current, the driver, the flyback diode in entry 9.5, and the enclosure volume for the whole life of the product.
- Mechanical relays do have their place, for high-current switching, galvanic isolation requirements, and certain safety certifications, but most of the time you do not actually need those properties.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
A MOSFET for DC loads	For DC loads, a MOSFET is often the first thing to consider: tiny, silent, essentially unlimited cycles. A logic-level part on the low side can be driven from a microcontroller pin with a gate resistor and a pull-down.
A solid-state relay where isolation is required	For AC loads or where the safety standard requires isolation, a solid-state relay gives you control isolation without the moving parts, but it leaks a little when off and tends to fail shorted, so it is not a safety disconnect.
A mechanical relay only when its properties are required	Before locking one in, ask whether you actually need the specific properties it provides, like true galvanic isolation or very high current.

BEFORE YOU COMMIT

- **Cycle count over product life:** Multiply the switching events per day by the product life in days and compare it with the relay's rated cycles at your load.
- **Coil current in the power budget:** If a relay stays, its coil current is in the power budget for the entire time it is energized.
- **Isolation requirement in writing:** The safety standard for your product category says whether isolation is required, and that is what decides between a MOSFET and a relay.

Simple to understand is not simple to ship

A relay is the easiest switch to explain and one of the hardest to live with, because it is the only part on the board that wears out mechanically every time it does its job.

STAGE: SCHEMATIC AND INDUSTRIAL DESIGN

CONNECTORS, INTERFACES, AND PERIPHERALS

The connectors and interfaces on a product are where the customer touches your engineering, and where a prototype habit becomes a return. Three of these entries are the connectors that fail mechanically or frustrate customers, and four are the interface and peripheral technologies that look modern on a spec sheet and turn into the hardest part of the project for a small team. Every one has a simpler option that gets you to production.

IN THIS PART

8.1	USB Micro-B Connectors	p. 74
8.2	Barrel Jack Power Connectors	p. 75
8.3	Bare 2.54 mm Pin Headers as Production Connectors	p. 76
8.4	Proprietary or High-Speed Connectors and Interfaces	p. 77
8.5	Advanced USB Features (USB 3.0, Power Delivery)	p. 78
8.6	Fancy Displays (AMOLED and MIPI DSI Panels)	p. 79
8.7	High-End Analog Audio	p. 80

8.1 USB Micro-B Connectors

CONNECTORS

WHY IT FAILS IN A PRODUCT

- I will be honest, I hate Micro-B connectors. I went through multiple Kindles over the years, and the Micro-B port was always the first thing to fail. That frustration is not just mine, it is a real engineering problem.
- USB Micro-B is the most notorious mechanical failure point in consumer electronics. The receptacle on your board holds a thin tongue that carries the contacts, and lateral stress on the plug snaps that tongue or shears the receptacle's surface-mount pads off the PCB.
- In practice it fails far sooner than USB-C does, whatever the cycle rating on the datasheet says, and EU rules already require USB-C for most categories of consumer electronics that charge over a cable, so a Micro-B product has a compliance problem in Europe before it has a reliability problem anywhere else.
- Customers increasingly expect USB-C on new products, and a Micro-B port signals that your product is outdated before it even ships.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
A USB-C receptacle	Reversible, far more tolerant of side loads, and the connector customers already have a drawer full of cables for.
A mid-mount or through-hole-anchored USB-C part	Choose a connector with through-hole shell anchors, so the mechanical stress goes into the board and not into the signal pads.
The 5.1 kOhm resistors on CC1 and CC2	A USB-C power sink needs a pull-down on each CC pin, or a USB-C to USB-C cable will never deliver power to it, which is the most common USB-C mistake I see.

BEFORE YOU COMMIT

- **Shell anchoring:** The connector has through-hole shell tabs, and the enclosure supports the connector body so the cable cannot lever it off the board.
- **CC pull-downs present:** Both CC pins have 5.1 kOhm pull-downs to ground, and the product has been tested with a C-to-C cable from a USB-C charger.
- **EU common charger rule:** If the product sells in the EU and charges over a cable, confirm it falls under the USB-C requirement and the labeling that comes with it.

There is almost no reason left

For a product designed in 2026, there is almost no scenario where starting with USB Micro-B is the right call, and the EU has made that decision for you on most consumer products.

8.2

Barrel Jack Power Connectors

CONNECTORS

WHY IT FAILS IN A PRODUCT

- Barrel jacks are frustrating from an engineering standpoint, and even worse from a customer experience standpoint. With USB-C, your customer already has a dozen chargers lying around the house that will work. With a barrel jack, they have to keep track of the one adapter that came in the box.
- I have owned products where I lost the barrel jack charger, and finding a replacement with the right size, right polarity, and right current rating was such a hassle that I just bought a whole new product instead. That is not the experience you want your customers to have.
- On the engineering side, there are too many similar but incompatible sizes. A 5.5 by 2.1 mm plug looks almost identical to a 5.5 by 2.5 mm plug, but one will not fit right in the other's jack.
- There is no built-in polarity protection either, and no voltage negotiation, so a customer who plugs in the wrong adapter can fry the board instantly, as entries 9.1 and 9.4 explain.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
USB-C for low-to-moderate power	For most consumer products, USB-C is the better default, with 5 V at up to 3 A available when the charger advertises it on the CC line, which your design reads before drawing more than the USB default.
USB-C Power Delivery when you need more	If the product needs more than 15 W, USB PD can negotiate up to 100 W or more, but read entry 8.5 before you sign up for a PD controller.
A locking connector for high current	If you need more current than USB-C can handle, use a locking connector rated for your specific current, with a keyed shape and a supply you provide.

BEFORE YOU COMMIT

- **Power budget:** The product's peak power draw fits inside 15 W for plain USB-C, or you have planned the PD negotiation.
- **Wrong-adapter test:** If a barrel jack stays, the board survives a reversed-polarity adapter and a 12 V adapter on a 5 V input without damage, with the protection to prove it.
- **Replacement path:** The customer can buy a replacement power source without contacting you, which is only true when the connector is a standard.

The junk drawer test

Two years after purchase, the customer will grab whatever adapter fits. Design for that customer, not for the one who keeps the original adapter on a labeled hook.

8.3

Bare 2.54 mm Pin Headers as Production Connectors

CONNECTORS

WHY IT FAILS IN A PRODUCT

- Bare pin headers are perfect for prototyping, and they are terrible for anything that actually ships to a customer.
- They have no locking mechanism and no strain relief, which means they vibrate loose in anything that gets moved, shipped, or even handled roughly, and the first symptom is an intermittent fault nobody can reproduce.
- They are easy to plug in offset by one pin since there is no keying to prevent it, and that can damage your board or connected components instantly, on the assembly line or in a repair.
- The exposed pins also make emissions testing harder, because every unshielded wire hanging off a header is an antenna.

SWAP IT FOR

REPLACEMENT	WHY IT WORKS
JST PH, XH, or GH series	Keyed, latching wire-to-board connectors at pitches from 1.25 mm to 2.5 mm, which cover most internal cable connections in a consumer product.
Molex Pico-Blade or Micro-Lock	Compact keyed connectors for tight enclosures. Pico-Blade uses friction retention, and Micro-Lock adds a positive latch for higher-vibration products.
Hirose DF series	A good choice when the connection needs a high mating cycle count or a very low profile.

BEFORE YOU COMMIT

- **Retention and keying on every connector:** Every internal connection has a keyed housing and a latch or friction lock rated for the vibration the product will see.
- **Mating cycles:** The rated mating cycles cover assembly, test, rework, and any repair, since some small connectors are rated for only a few dozen.
- **Crimp tooling:** Your assembler has the crimp tool for the contact you chose, or you have specified pre-made cable assemblies.

Cost check

The cost increase over a bare header is usually a few cents per connection, and it is justified by the reliability and usability improvement every single time.

8.4

Proprietary or High-Speed Connectors and Interfaces

TECHNOLOGY TO AVOID

WHY IT FAILS IN A PRODUCT

- It is easy to be tempted by sleek, modern interfaces like MIPI, Thunderbolt, or Apple's Lightning connector. They look clean and offer blazing-fast data rates on paper, but integrating them is brutal.
- You will often need expensive test fixtures, detailed timing analysis, and access to documents you cannot get unless you are a high-volume partner.
- With Lightning you are also dealing with licensing fees and per-unit royalties through Apple's MFi program, and with Thunderbolt you are dealing with certification and a controller chip you cannot buy without a relationship.
- Unless you are shipping millions of units, you will not have the support or the bargaining power to use these interfaces safely.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Use SPI and I2C for sensors and small peripherals	They are not flashy, but they are forgiving, universally supported, and they never block you from getting to production.
Use LVDS or parallel RGB for displays	Accessible interfaces that mainstream microcontrollers and processors support directly, as entry 8.6 covers.
Use HDMI or USB for external video and data	Off-the-shelf bridge chips and test equipment a small team can rent. HDMI needs an adopter agreement and a small per-unit royalty, which is routine next to MFi or Thunderbolt.

BEFORE YOU COMMIT

- **Specification access:** You can download the full interface specification and the chip datasheet without an NDA or a membership.
- **Test equipment:** You know what it costs to verify the interface, and it is equipment you can rent for a week rather than buy.
- **Licensing:** There are no per-unit royalties or certification programs attached to the interface, or you have priced them into the product.

Flashy is expensive

The interfaces that look most impressive on a spec sheet are the ones with the fewest documents, the most expensive tools, and the least support for a small team.

8.5

Advanced USB Features (USB 3.0, Power Delivery)

TECHNOLOGY TO AVOID

WHY IT FAILS IN A PRODUCT

- USB 3.0 SuperSpeed and USB Power Delivery sound great on paper. You get higher data rates and flexible power, but both come with serious implementation challenges.
- USB 3.0 runs at 5 gigabits per second or more. To get that working reliably you need carefully matched differential pairs, tight impedance control, and a multi-layer stackup, and debugging it needs high-bandwidth oscilloscopes and signal integrity setups most small teams do not have.
- USB Power Delivery lets devices negotiate up to 20 V and 5 A, and newer versions go higher, but it means adding a dedicated PD controller and a power path built for the higher voltage, plus negotiation firmware unless the controller negotiates on its own from stored settings. A bug in that negotiation can push dangerous voltage into a 5 V rail, as entry 9.4 explains.
- Unless your product truly needs high-speed data or more than 5 V from USB, you are better off keeping things simple.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
USB 2.0 for data	High speed signals at 480 megabits per second and full speed at 12, and either covers firmware updates, logging, and most sensor data. Check which one your microcontroller actually has, since many only do full speed.
5 V USB-C charging for power	Plain USB-C with the CC pull-downs gives you up to 15 W with no controller chip and no negotiation firmware.
Power Delivery only when you need it	PD is great when the product needs more than 15 W, but you definitely want to know what you are signing up for, including the protection on the input.

BEFORE YOU COMMIT

- **Data rate, in numbers:** Write down the sustained data rate the product actually moves over USB. Full speed delivers a few megabits per second in practice and high speed a couple of hundred, so match the port.
- **Power, in numbers:** Write down the peak power the product draws from the port. If it is under 15 W, you do not need PD.
- **Input protection if PD stays:** The 5 V rail is protected against a PD negotiation gone wrong, with an overvoltage cutoff rated for the highest voltage the cable can carry.

Know what you are signing up for

Power Delivery and SuperSpeed are both fine technologies. They are also each a project of their own, and most products can ship without either one.

8.6

Fancy Displays (AMOLED and MIPI DSI Panels)

TECHNOLOGY TO AVOID

WHY IT FAILS IN A PRODUCT

- An active matrix OLED screen, like the kind used in smartphones, can make your product look incredibly sleek and modern. Despite how cool they look, these displays are usually a bad choice for small teams.
- They often require a MIPI DSI interface, which is not supported by basic microcontrollers, so you are locked into a high-end microcontroller with a DSI host or an application processor even if the rest of the product does not need that kind of power.
- The panels themselves are priced for large-scale phone manufacturers, not low-volume startups, and they have short lifespans in the supply chain. A model might be discontinued without warning or replaced by a new revision that is not compatible with your existing design.
- Even when you can get documentation, integrating them takes real time: initialization sequences, power rails, brightness control, and timing all have to be worked out from thin documents.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
A standard TFT LCD	Suppliers like Newhaven or BuyDisplay.com sell panels with real datasheets, stable part numbers, and reasonable minimum orders.
Common resolutions and simple interfaces	Stick to common resolutions and SPI or parallel RGB interfaces, so a mainstream microcontroller or a low-cost processor drives the panel directly.
A panel with a published lifecycle	Choose a display vendor who commits to availability and offers a compatible replacement when a panel changes.

BEFORE YOU COMMIT

- **Interface supported by your processor:** The display interface is one your chosen chip drives natively, without a bridge chip or a processor upgrade.
- **Lifecycle and MOQ in writing:** The vendor has confirmed availability, lead time, and minimum order for the exact panel part number.
- **Documentation complete:** You have the initialization sequence, the timing, and the brightness control specification, plus the backlight requirements if it is an LCD, before you commit to the mechanical design around the panel.

A display decides the processor

The display interface is often what forces a product from a microcontroller to an application processor. Choose the panel with that consequence in mind, not from a photo.

8.7 High-End Analog Audio

TECHNOLOGY TO AVOID

WHY IT FAILS IN A PRODUCT

- Audio seems simple, it is just sound, right? High-end analog audio is one of the most unforgiving parts of electronics.
- Your layout matters, grounding matters, shielding matters, and power supply noise matters. Even tiny mistakes create hiss, hum, or distortion that customers will definitely hear.
- Validating those designs requires audio analyzers and measurement setups most small companies do not have access to, so you cannot even quantify the problem you are hearing.
- If you do not have serious analog expertise on your team, do not try to go toe to toe with audiophile-grade products.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Use digital audio interfaces	USB audio, I2S to a codec, or a Bluetooth audio module keep the signal digital until the last possible stage, where the hard work is handled for you.
Use an integrated codec or class D amplifier	Parts from TI, Cirrus, or Analog Devices put the sensitive analog stages inside one chip with a reference layout you can copy.
If you need analog, copy the reference design exactly	Start with something small and reproduce the vendor's reference layout, grounding, and supply filtering line for line.

BEFORE YOU COMMIT

- **Where does the signal go analog:** Draw the signal chain and confirm the analog section is as short as possible and lives inside a chip with a reference design.
- **Supply and ground plan:** The audio stages have their own filtered supply, and they are placed so the regulator's and the radio's return currents do not flow through the audio return path on the shared plane. Do not split the ground plane to get there.
- **Listen and measure:** Test the product in silence at full gain, on the real enclosure, with the radio transmitting, and measure the noise floor rather than trusting your ears alone.

Customers hear everything

A hum at the noise floor is invisible on a spec sheet and impossible to ignore in a quiet room. Analog audio is the one place where a small layout mistake is a one-star review.

STAGE: SCHEMATIC DESIGN

CIRCUIT PROTECTION YOU NEVER SKIP

If you learned electronics through kits and tutorials, your prototype is probably missing every protection in this part. Your prototype is your baby, and you treat it that way. You always use the right power adapter, and you would never yank it out mid-charge or cut the power in the middle of a firmware update. Your customers will not be that gentle. They will plug in whatever adapter fits, zap it with static, and cut the power at the worst possible moment. Two of these protections are all that stands between your product and an actual fire. The seventh, the watchdog and brown-out protection built into your microcontroller, lives in Part 11 with the firmware it belongs to.

IN THIS PART

9.1	Shipping Without Reverse Polarity Protection	p. 82
9.2	Shipping Without ESD Protection	p. 83
9.3	Shipping Without Overcurrent Protection	p. 84
9.4	Shipping Without Overvoltage and Transient Protection	p. 85
9.5	Shipping Without Inductive Kickback Protection	p. 86
9.6	Shipping Lithium Cells Without Dedicated Protection	p. 87

9.1 Shipping Without Reverse Polarity Protection

POWER INPUT

WHY PRODUCTS DIE WITHOUT IT

- A customer plugs in the wrong barrel jack adapter or inserts batteries backward, and one second later the board is dead. You would never make that mistake, but put a few thousand units into the market and someone absolutely will.
- The typical prototype fix is a series diode, and it works, but it drops 0.3 to 0.7 V across the diode, which wastes power and adds heat in a production product, and on a battery product it eats a meaningful slice of the usable voltage range.
- A keyed connector is not reverse polarity protection. Customers defeat mechanical safeguards all the time, and a battery holder with springs will accept a cell backwards no matter how it is labeled.
- Every input that a person can connect is a candidate: the DC jack, the battery contacts, and any screw terminal on the product.

ADD THIS

THE PART	HOW TO APPLY IT
A P-channel MOSFET on the high side	The MOSFET's body diode conducts under normal polarity, and with the gate tied to ground the resulting gate-to-source voltage turns the channel fully on, so instead of a diode drop you get almost no voltage drop at all.
An ideal diode controller for higher current	For higher-current designs, an ideal diode controller drives an N-channel MOSFET and does the same job even more efficiently.
A Zener and resistor on the gate at higher inputs	If the input voltage can exceed the MOSFET's gate rating, clamp the gate-to-source voltage with a Zener and a series resistor.

BEFORE YOU COMMIT

- **Gate voltage rating:** The maximum input voltage, including a wrong adapter, stays inside the MOSFET's gate-to-source rating, or the gate is clamped.
- **On-resistance and current:** The MOSFET's on-resistance at your peak current produces a voltage drop and a power dissipation you can live with.
- **Reverse test on a real unit:** Connect the supply backwards on a finished unit, then connect it correctly and confirm the product still works.

Cost check

Total cost for this protection is just a few cents, and it saves the first support call from the first customer who grabs the wrong adapter.

9.2 Shipping Without ESD Protection

EXTERNAL PORTS

WHY PRODUCTS DIE WITHOUT IT

- Your product survives all of your testing, then dies in a carpeted office in dry winter air. Every connector, button, and exposed port that a person can touch is an entry point for electrostatic discharge, and a simple carpet-to-doorknob shock can hit anywhere from 8 to 15 kV.
- This is not just a reliability concern, it is a certification problem. Selling in Europe means passing ESD immunity testing under IEC requirements, and if your product fails at the test lab, that means a board respin while your launch date slips.
- A lot of engineers assume the ESD structures built into their microcontroller pins are enough. Those internal structures are designed to survive the assembly line, not your customer's wool sweater in the middle of winter.
- A Hardware Academy member had a product working perfectly for a year, then a unit came back with the charger IC physically blown. His USB shield was grounded, but he had no ESD diode on the USB power pin, and one static discharge through the port killed the charger IC.

ADD THIS

THE PART	HOW TO APPLY IT
ESD diodes on every externally accessible line	Place them right at the connector with a short, direct return path to ground, so the spike is clamped before it can reach your sensitive components.
Low-capacitance parts on high-speed lines	For USB and other high-speed data lines, use low-capacitance versions so you do not degrade signal integrity.
Protection on power pins too	The USB power pin, the charger input, and every button line need a clamp as much as the data lines do.

BEFORE YOU COMMIT

- **Every external line listed:** Walk the enclosure and list every pin, pad, button, and shield a person can touch. Each one has a clamp or a documented reason it does not need one.
- **Placement and return path:** Each ESD diode sits between the connector and the rest of the circuit, with a ground via right at its pad, as entry 10.9 explains.
- **Pre-compliance zap test:** Test every port with an ESD gun before the real lab visit. Many pre-compliance labs rent time by the hour, and a rubbed-sweater test tells you almost nothing.

Cost check

Each protected line adds just a few cents, which is a fraction of the cost of one returned unit, let alone a board respin after a failed immunity test.

9.3

Shipping Without Overcurrent Protection

SHORT CIRCUIT

WHY PRODUCTS DIE WITHOUT IT

- A damaged cable, a shorted component, or a curious customer poking something where they should not, any of these can create a short circuit on your board. Without overcurrent protection, that short becomes smoke, melted PCB traces, and a safety incident with your company's name attached to it.
- This is where product risk turns into legal risk, because safety standards and major retailers expect overcurrent protection, and fire liability has ended small companies.
- When I was designing power management chips at TI, we built current limiting into every chip just to protect the silicon itself. Your board needs that same kind of protection at the system level, because without it one short takes out everything.
- A battery makes it worse, because a lithium cell can deliver tens of amps into a short, far more than any wall adapter, as entry 5.5 covers.

ADD THIS

THE PART	HOW TO APPLY IT
A resettable PTC fuse	For most consumer products this is the simplest answer at a few cents. A PTC sharply increases its resistance when too much current flows, tripping like a breaker and resetting itself once the fault clears and it cools.
An eFuse IC for faster, more precise limiting	PTCs are relatively slow, so they will not catch fast transients. An eFuse handles precise current limiting along with overvoltage lockout in the same part, typically for about a dollar.
A traditional one-time fuse on safety-critical rails	It blows permanently, which is exactly what you want when the alternative is a fire, and it is your last line of defense behind the others.

BEFORE YOU COMMIT

- **Hold and trip currents:** The PTC's hold current sits above your peak normal load at the highest ambient temperature, and its trip current is below what the traces and connectors can survive.
- **Trace capacity:** Every trace and via between the source and the protection can carry the fault current long enough for the protection to act, as entry 10.3 covers.
- **Short test:** Short the output of a finished unit through a rated fixture with a current-limited source or a fused lead, never bare clip leads on a lithium pack, then remove the short and confirm the product recovers.

Layers, not a single part

A PTC for the everyday fault, an eFuse where precision matters, and a one-time fuse behind them on anything that could burn. Each one catches what the others miss.

9.4

Shipping Without Overvoltage and Transient Protection

POWER INPUT

WHY PRODUCTS DIE WITHOUT IT

- You ship your product with a 5 V adapter. Two years later the customer grabs a 12 V adapter from the junk drawer, and it fits because it is the same barrel jack. Everything downstream sees too much voltage.
- A wrong adapter is not the only way an overvoltage condition occurs. Voltage spikes from hot-plugging cables under load, line transients from noisy power sources, and load-dump surges if your product goes anywhere near a vehicle electrical system all cause the same kind of damage.
- A TVS diode can only absorb short bursts of energy, so it cannot protect you from a wrong adapter feeding your board 12 V all day long, and a sustained overvoltage overheats the TVS, which usually fails short and then burns open, taking the board with it.
- If your product takes power through USB-C, USB Power Delivery can negotiate up to 48 V on the same cable, so a misbehaving charger or a bug in your PD negotiation can push dangerous voltage into a 5 V rail.

ADD THIS

THE PART	HOW TO APPLY IT
A TVS diode on the power input	Size it so its standoff voltage sits above your maximum normal operating voltage but its clamping voltage stays below the damage threshold of your downstream parts.
An eFuse or overvoltage protection IC	For sustained overvoltage like a wrong adapter, you need a part that shuts the input off entirely, so the wrong adapter gets ignored instead of killing your board.
An OVP threshold set for the input you designed	On a 5 V-only USB-C input, set the cutoff just above 5 V with margin for tolerance, so a PD mishap trips before anything sees 9 or 20 V. A PD input needs a power path built for the negotiated voltage.

BEFORE YOU COMMIT

- **Standoff, clamp, and damage voltages:** Write down the TVS standoff, the TVS clamp at the surge current, and the lowest absolute maximum rating downstream. They have to be in that order.
- **Wrong-adapter test:** Feed a finished unit 12 V on a 5 V input from a current-limited supply. The product should refuse to run and then work normally on the right adapter.
- **Automotive:** If the product ever plugs into a vehicle, the input protection is rated for load dump, which is a much bigger event than a wrong adapter.

Cost check

All of that protection adds well under a dollar, and it is the difference between a customer shrugging at a dead adapter and a customer returning a dead product.

9.5

Shipping Without Inductive Kickback Protection

COILS AND MOTORS

WHY PRODUCTS DIE WITHOUT IT

- Anything with a coil in your product fights back when you cut the current. Relays, motors, and solenoids all store energy in a magnetic field while current flows, and the moment you switch that current off, the collapsing field fires a high-voltage spike backward into your driver circuits and microcontroller pins.
- The sneaky part is that this usually survives short bench tests just fine. The damage is cumulative, so units start failing in the field one at a time, weeks or even months apart, and it looks like random bad luck until you finally figure out the pattern.
- The spike also couples into nearby signals, so a product with an unprotected relay often has mystery resets and corrupted readings long before the driver transistor finally dies.
- Brushed DC motors and solenoids produce the same spike as a relay coil, and they are the ones most often left unprotected because nobody thinks of a motor as an inductor.

ADD THIS

THE PART	HOW TO APPLY IT
A flyback diode across every DC coil	Across every DC relay coil, solenoid coil, and one-direction DC motor, so the stored energy has a safe path to dissipate. A motor on an H-bridge uses the driver's own recirculation path instead.
Placed as close as possible to the coil	Not back at the driver, because the shorter that path is, the less chance the spike has to cause damage along the way.
A TVS or RC snubber for faster release or AC loads	A plain diode slows a relay's release time. Where that matters, or on an AC coil, use a TVS or an RC snubber sized for the coil.

BEFORE YOU COMMIT

- **Every coil found:** Walk the schematic and list every relay, solenoid, motor, and transformer winding, and confirm each has a flyback path.
- **Diode rating:** The diode's current rating covers the coil current and its reverse voltage rating covers the supply, with margin.
- **Scope the switch node:** Watch the driver transistor's collector or drain on a scope while the coil switches off, and confirm the spike is clamped below the transistor's rating.

Cumulative damage looks like bad luck

A transistor that takes a small overvoltage spike a thousand times a day does not fail on the bench. It fails in month four, one unit at a time, and every return looks unrelated to the last.

9.6

Shipping Lithium Cells Without Dedicated Protection

BATTERY

WHY PRODUCTS DIE WITHOUT IT

- If your product has lithium batteries and no dedicated protection circuitry, you are setting yourself up for a recall.
- Over-discharging a cell drops its voltage below the point of safe recovery and gradually kills the pack, sometimes without any visible sign until the battery will not take a charge at all. Overcharging and short circuits cause swelling, venting, and fire.
- You need UN38.3 testing just to legally ship lithium cells, plus additional safety certifications depending on your market, and most large retailers stack their own requirements on top of that, as Part 5 covers in detail.
- This is the second item in this part that can literally burn a house down, and the one where your own charging circuit might not be fully working yet during development.

ADD THIS

THE PART	HOW TO APPLY IT
Protected cells first	Start with cells that already have the safety circuit built in. That gives you overvoltage, undervoltage, overcurrent, and short-circuit protection right at the cell, it reduces your liability, and it keeps you safer during development.
A dedicated charger IC with correct termination	On your board, add a charger IC with proper charge termination and temperature sensing, so cells are not charged above their rated limit or outside their temperature window.
A thermistor on the cell	Wire the charger's temperature input to a thermistor on the cell, so charging stops outside the cell's charging temperature window.

BEFORE YOU COMMIT

- **Protection on the shipped cell:** The cell or pack you will actually ship has a protection circuit with documented thresholds, not just the sample you tested.
- **Charge termination verified:** On the bench, charging terminates at the cell's rated voltage and taper current, and charging stops outside the cell's temperature window rather than on any warm cell.
- **Documentation:** The UN38.3 summary and the cell datasheet are on file, as entry 5.7 requires.

Two layers, always

Protection at the cell and a proper charger on the board are not alternatives. The cell protection catches the fault your charger causes, and the charger keeps the cell from ever needing it.

STAGE : PCB LAYOUT

PCB LAYOUT RULES YOU NEVER BREAK

Two boards can be built from the exact same schematic, and one works perfectly while the other fails EMI testing, resets randomly, or has almost no wireless range. The schematic is really just an approximation of what actually gets built. The layout is where the real physics happens, and small placement and routing decisions can make or break your entire product. These rules run in the order you meet them in a layout: power and placement first, then signals, then RF, then the test points that let a factory build the board at all.

IN THIS PART

10.1	Never Place Decoupling Caps Far From the Pin	p. 89
10.2	Never Ignore the Switching Regulator Hot Loop	p. 90
10.3	Never Starve a Power Trace or Via	p. 91
10.4	Never Skip a Thermal Path for Hot Components	p. 92
10.5	Never Use Via-in-Pad Without Knowing the Cost	p. 93
10.6	Never Route Signals Across a Plane Split	p. 94
10.7	Never Route Differential Pairs Like Ordinary Traces	p. 95
10.8	Never Run High-Speed Traces Near the Board Edge	p. 96
10.9	Never Leave the Board Without Stitching Vias	p. 97
10.10	Never Put Copper Under the Antenna	p. 98
10.11	Never Design a Board Without Test Points	p. 99

SECTION

PART 10

DOC NO.

NVR-01

ENTRIES

11

SOURCE

[PREDICTABLEDESIGNS.COM](https://predictabledesigns.com)

10.1 Never Place Decoupling Caps Far From the Pin

POWER INTEGRITY

WHAT GOES WRONG WHEN YOU BREAK IT

- Decoupling capacitors only work when they are close to the power pin they are protecting. A decoupling cap acts like a tiny local energy tank that supplies the chip during sudden current spikes, faster than the main power supply can respond.
- Every millimeter of trace between the cap and the pin adds inductance, and inductance resists exactly those fast current changes. Put the cap too far away and you have built a filter that blocks the cap from doing its one job.
- The chip's supply dips during every burst of activity, and you get random glitches and resets that never correlate with anything in the firmware.
- A cap that is close but connected through a long thin trace to ground has the same problem, because the return path counts just as much as the supply path.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Place the smallest cap closest to the pin	The 100 nF cap goes right at the pin, with any larger bulk capacitor behind it, since the small one handles the fastest edges.
Connect it with short, wide traces	A short, wide trace from the pad to the pin and from the pad to ground, with no thermal relief on the connection.
Drop the ground vias right at the cap's pads	The via to the ground plane sits at the capacitor pad, not at the end of a trace, and two vias are better than one.

BEFORE YOU COMMIT

- **One cap per power pin:** Every power pin on every IC has its own capacitor within a few millimeters, placed before any signal routing starts.
- **Loop drawn:** For the fastest chips, trace the loop from the pin through the cap to ground and back on the layout, and confirm it is as small as physically possible.
- **Datasheet followed:** The values and placement match the chip vendor's decoupling recommendations, which are usually in the hardware design guide rather than the datasheet.

Place these first

Decoupling capacitors are the first parts to place after the chip itself, because once the signals are routed there is never room to move them closer.

10.2 Never Ignore the Switching Regulator Hot Loop

SWITCHING REGULATORS

WHAT GOES WRONG WHEN YOU BREAK IT

- Every switching regulator has a hot loop, which is the current path where current starts and stops in just a few nanoseconds. That fast-changing current makes the loop behave like a tiny transmitting antenna, and the bigger the loop area, the more noise it sprays across your board.
- Chip makers like TI and Analog Devices publish recommended layouts in their datasheets, and those layouts exist specifically to keep this loop tiny.
- A large hot loop shows up as radiated emissions at the test lab, as noise in every ADC reading, and as reduced sensitivity in any radio on the board.
- For a buck converter the hot loop is the input capacitor, the high-side switch, and the low-side switch or diode. For a boost it is the output capacitor and the switches. Know which one your part has.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Place the input capacitor as close as physically possible	The input cap for a buck sits right at the VIN and ground pins, before anything else around the regulator is placed.
Keep the loop on a single layer	Route the hot loop on the same layer as the part with no vias in it, and put a solid ground plane directly underneath.
Copy the datasheet layout instead of improvising	The vendor already solved this. Reproduce their placement and their copper shapes, including which pins get the ground via.

BEFORE YOU COMMIT

- **Loop identified:** You have drawn the hot loop for your specific topology on the layout and measured its area.
- **Inductor and switch node:** The switch node copper is as small as the current allows, and the inductor is shielded, as entry 6.10 covers.
- **Sensitive circuits away:** The analog front end, the antenna, and any audio path are on the far side of the board from the regulator.

The datasheet layout is the design

For a switching regulator, the reference layout is not a suggestion. The vendor's applications team already tested it, and any improvement you make is more likely to be a regression.

10.3 Never Starve a Power Trace or Via

CURRENT CAPACITY

WHAT GOES WRONG WHEN YOU BREAK IT

- A power trace that is too thin drops voltage under load, heats up, and in the worst case burns right off the board. The same goes for vias, since a single small via might only handle an amp or so before it becomes the bottleneck.
- Product creators often route power using the same default trace width as their signals, and everything seems fine at low current during early testing.
- Then the product hits a high-load condition, the supply rail sags, and the microcontroller resets at the worst possible moment, which looks exactly like a firmware bug.
- A starved trace also defeats the overcurrent protection in entry 9.3, because the trace can fail before the fuse or the PTC ever trips.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Size every power trace with a calculator	Trace width calculators are free and take seconds to use, so size every power trace for its actual current with plenty of margin and a temperature rise you accept.
Use copper pours for higher currents	Above an amp or two, a pour is easier to route and carries far more current than a trace, on any layer.
Place multiple vias in parallel wherever power changes layers	Several vias share the current and keep the connection alive if one is poorly plated.

BEFORE YOU COMMIT

- **Peak current per rail:** You know the peak current on every rail, including motor starts and radio bursts, and every trace on that rail is sized for it.
- **Via count at every layer change:** Power that changes layers goes through at least two vias, sized from the same calculator.
- **Voltage drop under load:** Measure the rail at the load on a finished board at peak current, and confirm the drop is inside the regulator's and the chip's tolerance.

Copper is free

Wider traces and bigger pours cost nothing on a board you are already paying for. There is no reason to route power at the minimum width your tool defaults to.

10.4 Never Skip a Thermal Path for Hot Components

THERMAL

WHAT GOES WRONG WHEN YOU BREAK IT

- A voltage regulator that works fine for 10 minutes and then browns out is commonly a layout problem, not a component problem.
- Many regulators and other hot components dissipate their heat through an exposed pad on the bottom of the package called a power pad. That pad needs somewhere for the heat to go, and on most boards that means an array of thermal vias connecting it down to a large copper pour on another layer.
- Skip those vias and the heat just piles up until the part hits thermal shutdown, and then your product resets over and over out in the field, worst of all on a hot day inside a sealed enclosure.
- Those thermal vias are technically via-in-pad, but a power pad is the one place where it is normally acceptable, since a little solder wicking there mostly causes small voids instead of failed joints.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Follow the land pattern in the datasheet exactly	Including the recommended via array, the via size, and the paste mask pattern that keeps the part from floating during reflow.
Give the part as much copper area as you can spare	Copper is free, and it is the cheapest heatsink you will ever get, on the top layer and on the plane the vias connect to.
Calculate the dissipation before you place the part	Input-to-output voltage times current for a linear regulator, and the datasheet's thermal data, read with its test board conditions, tells you roughly how much copper you need.

BEFORE YOU COMMIT

- **Junction temperature estimated:** For every part with a power pad, you have estimated the junction temperature at the highest ambient inside the enclosure.
- **Via array present:** The power pad has the vendor's via array, connected to a copper pour, not just a pad with a single via.
- **Thermal camera on the first build:** Run a finished unit at full load in the enclosure and look at it with a thermal camera before you order production quantities, remembering the junction runs hotter than the case you can see.

The ten-minute failure

Thermal problems are the failures that never show up in a quick bench test, because the board is still cool. They show up after lunch on a hot day, in the field.

10.5 Never Use Via-in-Pad Without Knowing the Cost

ASSEMBLY AND FAB COST

WHAT GOES WRONG WHEN YOU BREAK IT

- Via-in-pad means placing a via directly inside a component's solder pad instead of next to it. It is tempting because it saves space, and for dense parts like fine-pitch BGAs it is sometimes the only option.
- The first cost shows up on the assembly line. During reflow the molten solder wicks down into the open via like a straw, which starves the joint and leaves you with weak or even open connections.
- The second cost shows up on every fab invoice for the life of the product, because the fix for that wicking is having your board fab fill each via with epoxy and plate copper over the top, and that process adds real money to every single board you ever build.
- An open joint from wicking is intermittent, so the boards pass a quick test and fail in the field when they warm up or flex.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Place vias next to the pad	Unless you truly need it, put the via beside the pad and connect it with a short trace, which costs nothing at the fab.
Specify filled and capped vias if you must	If your design does require via-in-pad, specify filled and capped vias in your fab notes, and budget for that extra cost from day one.
Reserve it for the parts that need it	Fine-pitch BGAs and power pads are where via-in-pad earns its place, and nowhere else on the board.

BEFORE YOU COMMIT

- **Via-in-pad count:** Run a check for vias inside SMD pads and justify each one, or move it.
- **Fab notes:** If any remain outside a power pad, the fab drawing calls out filled and capped vias explicitly.
- **Quote with and without:** Get a board quote with and without the filled-via process, so the cost of the convenience is a number on every unit.

Two costs most people never see coming

The first cost is a weak joint on the assembly line. The second is on every fab invoice for the life of the product. Know both before you place a single via inside a pad.

10.6 Never Route Signals Across a Plane Split

RETURN PATHS

WHAT GOES WRONG WHEN YOU BREAK IT

- Routing a signal across a split in your ground or power plane is the most invisible and most expensive mistake in this part. Every signal current has to return to its source, and at high frequencies that return current flows in the plane directly underneath the trace.
- When the trace crosses a split, the return current cannot follow it, so it detours all the way around the gap. That detour creates a huge loop that radiates noise, couples interference into every nearby circuit, and is a leading cause of both EMI failures and mystery glitches nobody can reproduce.
- Nothing about it looks wrong, since the trace appears perfectly normal and every net checks out in your design software.
- Plane splits themselves are sometimes the right choice, like separating a sensitive analog section. The rule is about what you route over them.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Check every trace against the planes beneath it	Before you order boards, look at every layer with the plane layers visible, and find every trace that crosses a gap, including the gaps around vias and connectors.
Reroute anything that crosses a gap	Move the trace so its return path stays under it, or close the gap under that trace.
Bridge a necessary split with a stitching capacitor	If a signal must cross between two planes at different potentials, place a capacitor next to the crossing to give the return current a path, but never across a safety isolation barrier, where only a rated Y capacitor belongs.

BEFORE YOU COMMIT

- **Plane layers reviewed with signals:** You have reviewed the layout with only the plane layers and the signal layer above them visible, one layer at a time.
- **Antipad clusters:** Rows of vias and connector pins have not created an accidental slot in the plane under a signal.
- **Split justified:** Every intentional plane split has a written reason and a defined bridge point.

The most invisible mistake

Every other layout error leaves a clue on the screen. This one looks perfect, passes every design rule check, and fails at the test lab.

10.7

Never Route Differential Pairs Like Ordinary Traces

HIGH-SPEED SIGNALS

WHAT GOES WRONG WHEN YOU BREAK IT

- Differential pairs like USB, Ethernet, and HDMI carry equal and opposite signals on two traces, and the receiver only reads the difference between them. This makes them extremely immune to noise, because interference hits both traces equally and cancels out, but only when the layout is done correctly.
- For that to work, both traces need matched lengths, a controlled impedance, an unbroken reference plane underneath, and no stubs hanging off the lines.
- Get any of those wrong and the signals arrive skewed or reflected, which turns into random errors that are brutal to debug, like a USB device that enumerates on one computer and not another.
- A stub is anything that branches off the pair, like a via barrel that continues to an unused layer, and it reflects part of the signal back toward the source. Even a test pad sitting in the line adds capacitance that disturbs the impedance.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Use your software's differential pair tools	Your PCB software has built-in tools for routing differential pairs with matched lengths and controlled spacing, so use them for every pair.
Ask your fab for an impedance-controlled stackup	They will gladly tell you the exact trace width and spacing to use for 90 ohm USB or 100 ohm Ethernet on their materials.
Keep the reference plane solid and the stubs off	Route the pair over one unbroken plane, and put ESD protection in line at the connector rather than as a branch.

BEFORE YOU COMMIT

- **Impedance target per pair:** Every pair has its target impedance written on the fab drawing, with the stackup the fab confirmed.
- **Length match reported:** The design tool reports the length mismatch on every pair, and it is inside the interface specification.
- **Plane under the pair checked:** The pair never crosses a plane split, as entry 10.6 covers.

The fab already knows the numbers

You do not have to calculate impedance yourself. Tell the fab the target, and they hand you the width and spacing for their stackup, because they do it every day.

10.8 Never Run High-Speed Traces Near the Board Edge

EMI

WHAT GOES WRONG WHEN YOU BREAK IT

- High-speed traces routed near the edge of the board are a classic way to fail EMI testing.
- Near the edge, the electric fields around a trace are no longer fully contained between the trace and its ground plane. Part of the field fringes out past the edge of the board, and that fringing field radiates like a little antenna.
- Everything will look fine on your bench, but the antennas at the test lab will pick it up immediately, and every failed EMI scan costs you money and weeks of schedule.
- The board edge is also where connectors and cables live, and a cable picks up that fringing field and carries it out of the enclosure as a much better antenna.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Keep fast signals away from the edge	At least 4 to 5 times the distance between the trace and the ground plane below it, which on most boards works out to around 1 mm or more.
Route them on inner layers when you can	A trace between two planes has its field contained on both sides, which is the quietest place on the board for a clock or a bus.
Fence the edge with stitching vias	A row of ground vias along the edge, as entry 10.9 covers, helps contain whatever field does reach it.

BEFORE YOU COMMIT

- **Edge clearance rule set:** Your design rules include a keepout for high-speed nets near the board outline, at the distance calculated from your stackup.
- **Clocks and buses located:** Every clock, USB pair, memory bus, and switching node has been checked for edge proximity on every layer.
- **Cables considered:** No fast trace runs parallel to a cable connector along the edge.

Money and weeks

A failed EMI scan is not just the lab fee. It is a board respin, a new assembly run, and a second lab visit, which puts the launch weeks behind for a trace that could have moved a millimeter.

10.9 Never Leave the Board Without Stitching Vias

GROUNDING

WHAT GOES WRONG WHEN YOU BREAK IT

- Stitching vias are the rows of ground vias that tie your ground planes together across the board.
- Without them, ground currents on different layers cannot take the shortest path back to their source, and the gaps between your planes start acting like slots that leak radiation.
- Stitching vias along the board edges also form a fence that helps contain noise inside the board.
- I recently reviewed a board that looked clean everywhere else but had almost no stitching vias, and we caught it right before the fab order went out. Adding them cost nothing, and it meaningfully improved that product's odds of passing FCC testing on the first try.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Sprinkle ground stitching vias generously	Across the whole board, on a rough grid, wherever there is ground copper on more than one layer.
Fence the edges	A row of vias along the board edge, spaced closely, ties the top and bottom ground together where the fields try to escape.
Surround high-speed and RF routing	Add vias around any high-speed routing, around RF circuits, and along antenna feedlines, where a return path matters most.

BEFORE YOU COMMIT

- **Via spacing:** Stitching vias are spaced no farther apart than a small fraction of the wavelength of your highest frequency, which for most consumer boards means a few millimeters near fast signals.
- **Every layer change has a return via:** Where a fast signal changes layers, a ground via sits right next to the signal via so the return current can follow.
- **Edges fenced:** The board outline has a ring of ground vias, especially near connectors and the antenna.

Free, and easy to forget

Stitching vias cost nothing, take minutes to add, and are missing from more boards I review than any other single item on this list.

10.10 Never Put Copper Under the Antenna

RF

WHAT GOES WRONG WHEN YOU BREAK IT

- Copper placed under an antenna will kill your wireless range. An antenna works by coupling energy into the space around it, and a ground plane or any copper underneath detunes the antenna and absorbs the energy before it ever leaves the board.
- Every PCB or chip antenna and every pre-certified wireless module has a keepout zone defined in its datasheet, an area where you are not allowed to put copper, traces, or components on any layer. The same goes for metal mounting screws and standoffs.
- This one also has a certification angle, because a pre-certified module is only certified when it is used the way the datasheet specifies. Violate the keepout, and the module in your product no longer performs as certified, which can come back to bite you during compliance testing.
- The enclosure counts too. A battery, a metal bracket, or a cable routed over the antenna detunes it just like copper on the board.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Copy the dev board keepout exactly	Look at any dev board for a wireless module and you will find a bare region around the antenna. Reproduce that keepout exactly in your own layout. A few antenna types, like patch antennas, need ground beneath them instead, so the datasheet always wins over the rule of thumb.
Put the antenna at the board edge	Place the module so its antenna hangs at the edge or a corner of the board, pointing away from the rest of the electronics.
Keep the enclosure clear too	No metal hardware, battery, or cable within the keepout in three dimensions, and no metallic paint or coating on the plastic above it.

BEFORE YOU COMMIT

- **Keepout drawn on every layer:** The keepout is a rule area in the design, and the design rule check confirms nothing is inside it on any layer.
- **Mechanical review:** The enclosure model has been checked for screws, standoffs, batteries, and cables inside the keepout volume.
- **Range tested in the enclosure:** Wireless range is measured on a finished unit inside the real enclosure, not on a bare board.

Certified only as specified

The module's FCC grant assumes the keepout in the datasheet. Fill it with copper and you have a different radio, one nobody certified.

10.11 Never Design a Board Without Test Points

PRODUCTION TEST

WHAT GOES WRONG WHEN YOU BREAK IT

- A board with no test points cannot be tested in production. That means you are either shipping units that were never verified, or you are paying someone to probe every single board by hand.
- Factories test boards using a fixture with spring-loaded pogo pins, and those pins need dedicated copper pads to land on. If you did not put those pads in your layout, there is nowhere for the pins to go, and adding them later means a new board revision.
- Test points are also a lifesaver during debugging, especially with leadless packages like QFNs and BGAs where the pins are hidden underneath the chip and there is nothing to probe.
- The programming signals matter most, because the factory has to load your firmware onto every board, and the debug port has to be reachable to lock it afterward, as entry 11.6 explains.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Add test points to every power rail	Small exposed copper circles on each rail let the fixture confirm the regulators before anything else is tested.
Cover every communication bus and debug signal	I2C, SPI, UART, the programming lines, and the reset line, so the fixture can program, talk to, and lock every board. Keep pads off high-speed pairs, where a connector or boundary scan does the job instead.
Put them on one side, on a grid	Keep test points on the bottom side where possible, at a pitch and diameter your test fixture builder specifies.

BEFORE YOU COMMIT

- **Fixture builder consulted:** You have asked your assembler or fixture vendor for their test point size, spacing, and edge clearance rules before placing them.
- **Programming path defined:** The production programming and lock sequence runs entirely through test points or a connector the factory can reach.
- **Test plan written:** A one-page test plan lists what the fixture measures on each point and the pass limits.

Basically free

Test points are just small exposed copper circles placed on your signals and power rails. They cost nothing, and without them a factory cannot build your board.

STAGE : FIRMWARE

FIRMWARE MISTAKES YOU NEVER SHIP

The firmware mistakes in this part usually go unnoticed in a prototype, and none of them cause a problem until the product is already in customers' hands. One of them wears out the flash memory in a microcontroller in days instead of years. Another lets anyone copy your entire firmware. A third is now against the law in the UK and California and will block your CE mark in Europe. Several of the fixes are only a few lines of code, and two of them, the watchdog and brown-out detection, are the only protections on the whole product that cost nothing.

IN THIS PART

11.1	Debug Logging Left in the Shipping Build	p. 101
11.2	Blocking Delays in the Main Loop	p. 102
11.3	Hammering Flash or EEPROM with Constant Writes	p. 103
11.4	Watchdog Timer Never Enabled	p. 104
11.5	Brown-Out Detection Turned Off	p. 105
11.6	Debug Port Left Open with No Readout Protection	p. 106
11.7	Secrets Hardcoded in the Firmware Image	p. 107
11.8	No Secure, Recoverable Firmware Update Path	p. 108
11.9	One Default Password Shared Across Every Unit	p. 109

11.1 Debug Logging Left in the Shipping Build

SERIAL LOGGING

WHY IT FAILS IN A PRODUCT

- Almost every firmware project starts with print statements sprayed all over the code, and that is fine while you are developing. The problem is leaving all of it turned on in the build you actually ship.
- Each print call has to format a string and push it out the serial port one character at a time, and with the default serial setup on most microcontrollers that stalls the rest of your main code while it runs. A single long line can tie up the processor for a couple of milliseconds, long enough to miss a sensor sample or throw off a timing loop, and the chip cannot sleep while it is printing.
- The bigger problem is what you are printing. Development logs tend to include internal state, memory addresses, and sometimes the actual credentials the device is using, and anyone with a five dollar USB to serial adapter can clip onto that UART pin and read it.
- Production logging at a sane level is fine, and it is useful for debugging units that come back from the field.

BEFORE YOU SHIP

THE CHECK	WHY IT MATTERS
Put logging behind levels	The verbose output is compiled out of the release build, so it costs nothing at runtime, and only warnings and errors make it into the product.
Audit every remaining log line	Nothing that reaches the UART exposes internal state, memory addresses, or credentials.
Make the logging non-blocking if it stays	Use a buffered, interrupt-driven or DMA serial driver that drops or rate-limits low-priority lines when the buffer is full, so a log line never stalls the main loop.

BEFORE YOU COMMIT

- **Release build inspected:** Read the UART output of the actual release build for a full boot and a typical session, and confirm nothing sensitive appears.
- **Timing measured:** Confirm the longest log call in the release build cannot delay any time-critical task.
- **Level switchable in the field:** Decide whether support can raise the log level on a returned unit without a new firmware image, and how that access is protected.

A few lines of code

This is one of the cheapest fixes in the handbook. Put your logging behind levels and let the release build compile the verbose output away.

11.2 Blocking Delays in the Main Loop

TIMING AND POWER

WHY IT FAILS IN A PRODUCT

- The delay function is the first thing everyone learns, and it is the easiest way to make an LED blink. The problem starts when delay becomes the timing architecture for the whole product.
- Without an operating system, a blocking delay means your main loop sits and spins doing nothing for that entire time, so it cannot check a button, service a sensor, or respond to a wireless packet.
- A quick button press that happens entirely during a 500 ms delay is never detected, and you may never catch it in your own testing, because your presses may just happen to land outside of that delay. Once thousands of customers are pushing the button, it is not a rare event.
- It is also brutal for battery life, because a processor spinning in a delay loop draws full run current the whole time instead of sleeping.

BEFORE YOU SHIP

THE CHECK	WHY IT MATTERS
Move to a non-blocking design	A hardware timer or a simple state machine decides when things happen, and the processor sleeps in between.
Check the clock before acting	Even a basic loop that checks the clock before doing anything is a big step up from a blocking delay.
Use a small RTOS if the product is complex	A real-time operating system handles the scheduling once there are enough tasks to juggle, and most vendor SDKs ship with one.
Keep any delay to microseconds	A few microseconds to satisfy a peripheral's timing requirement is fine, and that is the only place a delay belongs.

BEFORE YOU COMMIT

- **Search the code for delay calls:** Every remaining busy-wait delay is a few microseconds inside a driver, with a comment saying why. Under an RTOS, a task delay that yields the processor is fine.
- **Worst-case loop time measured:** Toggle a pin at the top of the main loop and measure the longest period on a scope, then compare it with the shortest event the product must catch.
- **Sleep current measured:** The product spends most of its time in a sleep mode, and the measured average current proves it.

The bug you will never see on your bench

A press that lands inside a delay is not a rare event once thousands of customers are pushing the button, even if it never once happened while you were testing.

11.3

Hammering Flash or EEPROM with Constant Writes

NON-VOLATILE MEMORY

WHY IT FAILS IN A PRODUCT

- Flash memory and EEPROM both wear out, and every part has a finite write budget. A typical microcontroller's internal flash is rated for around 10,000 write and erase cycles per sector, and a dedicated EEPROM chip is usually rated somewhere around a million.
- Those sound like big numbers until you do the math on firmware that saves a setting every second, because at that rate a million cycles are gone in under two weeks.
- It usually looks innocent, like saving the volume level or a run-time counter to non-volatile memory every time through the loop.
- Once a cell wears out, reads start returning garbage and settings reset themselves, and there is no fixing it in the field because the memory is physically damaged.

BEFORE YOU SHIP

THE CHECK	WHY IT MATTERS
Write only when a value changes	Then wait a few seconds before committing, so a customer spinning a knob does not generate a hundred writes.
Spread continuous logging across the memory	Do not write to the same spot every time, so no single section wears out early. Most chip vendors offer a storage library that handles this wear leveling for you.
Use FRAM for heavy write loads	FRAM is a type of non-volatile memory that can handle practically unlimited writes, and entry 1.5 covers the microcontrollers that include it.

BEFORE YOU COMMIT

- **Write budget calculated:** Multiply every write path's worst-case rate by the product life and compare the total with the datasheet endurance, with a margin of at least ten times.
- **Power-loss safety:** A write interrupted by a power loss leaves the old value intact, which usually means a two-copy scheme with a checksum, with the copies in separate erase sectors.
- **Counted on the bench:** Log the number of writes during a typical day of use on a real unit, since the estimate is always lower than the measurement.

Do the math

One write per second is 86,400 writes a day. Against a million-cycle budget, that is less than twelve days of use before the part starts failing.

11.4 Watchdog Timer Never Enabled

FAULT RECOVERY

WHY IT FAILS IN A PRODUCT

- The watchdog timer is built into just about every microcontroller on the market, and unfortunately most products ship with it turned off, usually because it was disabled during debugging and never turned back on.
- A watchdog is a simple countdown that resets the chip unless your firmware keeps resetting the counter. If your code hangs, the counter stops getting reset, the timer runs out, and the chip reboots itself.
- Firmware hangs for all kinds of reasons that never show up in your own testing, like an I2C sensor holding the data line low or a rare buffer overflow that only happens once a week. Without a watchdog, the customer sees a product that just stops working, and most of them will return it rather than power cycle it.
- The mistake I see even when people turn it on is resetting it from a timer interrupt. That interrupt keeps firing even when the main loop is completely stuck, so the watchdog never trips and you have gained nothing.

BEFORE YOU SHIP

THE CHECK	WHY IT MATTERS
Enable the independent watchdog	Use the watchdog that runs from its own clock, so it keeps counting even if the main clock or the firmware fails.
Reset it from the main loop, never from an interrupt	Feed it only after you have confirmed every critical task is still running, so a stuck task cannot hide behind a healthy interrupt.
Record the reset reason	Read the reset cause register at boot and log it, so a watchdog reset in the field is a data point instead of a mystery.

BEFORE YOU COMMIT

- **Hang test:** Deliberately hang the main loop on a test build and confirm the product recovers on its own within the watchdog period, and that the reset clears the cause, since a stuck I2C peripheral needs a bus-clear sequence or a power cycle, not just a reboot.
- **Interrupt audit:** No interrupt handler anywhere in the code touches the watchdog.
- **Timeout chosen:** The watchdog period is longer than the slowest legitimate operation, including any flash erase or wireless connection, and shorter than what a customer would notice.

Free, and already inside the chip

Some engineers argue that a watchdog just masks bugs, and I get where that comes from, but in the field a product that recovers itself beats one that stays locked up, every single time.

11.5 Brown-Out Detection Turned Off

POWER INTEGRITY

WHY IT FAILS IN A PRODUCT

- Brown-out detection is a circuit inside the microcontroller that watches the supply voltage and holds the chip in reset whenever the voltage drops below a safe level. When I was designing power management chips at TI, every chip had to have this kind of undervoltage lockout, because a chip running on too little supply voltage behaves unpredictably.
- A microcontroller with insufficient supply voltage does not just stop. It skips instructions and corrupts registers while it thinks everything is fine, and that can happen from a slow supply ramp at power-on just as easily as from a voltage dip.
- Now picture that in the middle of a flash write. A battery sags under load or a motor kicks on, the voltage dips for a few milliseconds, and the firmware is halfway through saving your settings. What ends up in flash is corrupted, and depending on what was being written, the unit becomes non-functional.
- A Hardware Academy member ran into a microcontroller that kept brown-out resetting, and the common reaction is to make the resets stop by turning off the detection. The advice he got was to be grateful the detector was tripping, because it was catching a real power problem that needed fixing.

BEFORE YOU SHIP

THE CHECK	WHY IT MATTERS
Keep brown-out detection enabled	If the detector keeps tripping, it is catching a real power problem. Turning it off only masks that problem.
Set the threshold above the flash-write minimum	The threshold belongs above the minimum voltage the chip needs for reliable flash writes, not just the minimum it needs to run.
Add a supervisor for a tighter threshold	An external voltage supervisor chip is a good option when you need precision, but never remove the built-in detector with nothing in its place.

BEFORE YOU COMMIT

- **Threshold versus flash spec:** The brown-out threshold in the option bytes or fuses is above the minimum flash programming voltage in the datasheet, with margin for the threshold's own tolerance.
- **Dip test:** Pull the supply down slowly and quickly on a finished unit while it writes to flash, and confirm the settings survive intact.
- **Root cause found:** Every brown-out reset seen during development has been traced to its cause, whether a starved trace, a weak regulator, or a sagging battery, and fixed.

The detector is the messenger

A brown-out reset is the chip telling you the power supply is not good enough. Disabling the detector does not fix the supply. It just removes the warning.

11.6 Debug Port Left Open with No Readout Protection

FIRMWARE PROTECTION

WHY IT FAILS IN A PRODUCT

- The debug port, usually SWD or JTAG, is how you program and debug the chip during development. Ship a product with that port wide open and anyone with a debugger can pull your entire firmware off the chip in seconds, and that is your whole product in a file.
- For a lot of products the hardware is the easier part to copy, especially if you are not using any custom chips, so your firmware is often the only thing standing between you and a knockoff.
- Most product creators never turn protection on, usually because nobody wants to lose the ability to debug a returned unit.
- An open debug port also exposes the secrets covered in entry 11.7, because reading the flash is the first step in finding them.

BEFORE YOU SHIP

THE CHECK	WHY IT MATTERS
Enable readout protection as the last programming step	Set it at the end of your production programming, after the firmware is loaded and tested, so you keep full debug access until the unit is finished.
Use your chip family's built-in setting	STM32 has readout protection levels, Nordic parts have an access port protection setting, and the ESP32 needs a permanent fuse that disables the debug port plus flash encryption to keep the firmware unreadable.
Pick a level you can still erase	The highest readout level on most STM32 parts can never be undone even by you, so choose the level that blocks readout but still allows a full erase for warranty repairs.

BEFORE YOU COMMIT

- **Read attempt on a locked unit:** Connect a debugger to a finished, locked unit and confirm the flash cannot be read.
- **Erase path confirmed:** Confirm you can still mass-erase and reprogram a locked unit for repair, before you lock a thousand of them.
- **Lock step in the production script:** The lock is a step in the automated programming sequence, not a manual action someone can skip.

One setting

Almost every chip family has this fix built in, and it is usually one setting in your production programming sequence.

11.7 Secrets Hardcoded in the Firmware Image

CREDENTIALS

WHY IT FAILS IN A PRODUCT

- Hardcoding private API keys, cloud secrets, or Wi-Fi credentials into the firmware is one of the most common mistakes, because it is the fastest way to get a prototype talking to the cloud.
- A firmware image is just a file, and every plain text string inside it is sitting there in the open. Anyone who pulls the image off the chip, or downloads it from your update server, can run one command that lists every readable string in the file.
- Since the same key is baked into every unit, one compromised device can compromise every device you have sold. That is when somebody starts pushing fake data into your cloud, running up your API bill, or pulling other customers' data.
- Wi-Fi credentials are a slightly different case, because the ones that end up hardcoded are usually your own, left over from prototyping on your office network, and now anyone who reads your firmware can get onto that network.

BEFORE YOU SHIP

THE CHECK	WHY IT MATTERS
Provision Wi-Fi through the customer	The customer's network details come in through a provisioning step, through your app or a temporary setup mode, so no Wi-Fi password ever lives in the firmware image.
Make cloud credentials unique per device	Program them at the factory or generate them on first boot, and store them in a secure element or a protected region of memory rather than in the main image.
Run the strings check on your own image	Before somebody else does, list every readable string in the release image and confirm no key, password, or URL you would not publish is in it.

BEFORE YOU COMMIT

- **Strings audit done:** The release image has been scanned for readable strings, and every hit is accounted for.
- **Per-device credentials:** Two units off the line have different cloud credentials, and revoking one does not affect the other.
- **Office network cleaned:** The Wi-Fi credentials used during development have been changed, whether or not they ever shipped.

One command

The tool that lists every readable string in a firmware image is free and takes a second to run. Run it on your own release build before anyone else does. It is one check, not the whole audit, since binary keys do not show up as text.

11.8 No Secure, Recoverable Firmware Update Path

FIELD UPDATES

WHY IT FAILS IN A PRODUCT

- Firmware bugs are about the only bugs you can fix after the product ships without getting units back from customers, but only if you built a way to do it. A surprising number of products ship with no update path at all, so the first serious bug you find becomes a permanent feature of every unit you have sold.
- The other version of this mistake is almost worse: an update path with no security on it, where the device accepts any firmware image somebody pushes to it. Anyone who figures out your update protocol can load their own code onto your customers' devices.
- An update that is not recoverable turns a failed download or a power loss halfway through into a non-functional unit.
- This has to be in place before the first unit leaves the factory, because you cannot add an update path to a unit that has no update path.

BEFORE YOU SHIP

THE CHECK	WHY IT MATTERS
Sign every firmware image	The device checks the signature before it runs any of the new code, using a public key stored in protected memory.
Keep two copies of the firmware in flash	The device runs the old copy until the new one is verified and then switches over, with a bootloader that confirms the new image actually boots before committing, so a failed update never leaves the unit dead.
Budget flash for two copies	You need room for both images, in internal flash or an external flash chip, so keep it in mind when you choose a microcontroller.
Set up the bootloader before the first unit ships	Most modern chip families come with a bootloader framework that handles this, but it has to be in place before the first unit leaves the factory.

BEFORE YOU COMMIT

- **Bad image rejected:** An image with a modified byte or a wrong signature is refused, and the unit keeps running the old firmware.
- **Power pulled mid-update:** Cut the power halfway through an update on a real unit and confirm it boots the previous image.
- **Rollback and version rules:** Decide whether older images can be installed, and how the device reports its version to the cloud.

Two things, before the first unit ships

A proper update path needs a signature check and a recoverable install. Both are far easier to build in on day one than to add after ten thousand units are in the field.

11.9 One Default Password Shared Across Every Unit

COMPLIANCE

WHY IT FAILS IN A PRODUCT

- This is the only firmware mistake in the handbook that is now against the law in parts of the world. Shipping every unit with the same default login was standard practice for years, and manufacturers sometimes even printed it in the manual so customers could get into their settings.
- Attackers know those default logins too, and there are automated tools that scan the internet for devices still using them.
- Since April 2024 the UK has banned universal default passwords on connected consumer products, with fines up to ten million pounds or 4 percent of global revenue. California has required unique or user-set passwords on connected devices that can be logged into remotely since 2020.
- In Europe, the cybersecurity rules for internet-connected wireless products that took effect in August 2025 treat a shared default password as a failure of the standard, which in practice means no CE mark through the normal self-certification route.

BEFORE YOU SHIP

THE CHECK	WHY IT MATTERS
Generate a unique password per unit	Create it at the factory and print it on the label, and store it in the protected memory covered in entry 11.7.
Or force the customer to set their own	Require a password change the first time the product is powered on, before any remote access is enabled.
Check every market you sell in	The UK rules, the California rules, and the European radio equipment cybersecurity requirements each have their own wording, and your compliance plan should cite all three.

BEFORE YOU COMMIT

- **Two units, two passwords:** Two units off the line have different passwords, and neither is derivable from the serial number or the MAC address.
- **Label process:** The factory prints the per-unit password on the label from the same record it programs into the unit, and the two cannot get out of sync.
- **Compliance file:** The password policy is documented in the technical file for the UK, California, and CE requirements.

Compliance warning

A shared default password is no longer just a security weakness. In the UK, California, and Europe it is now a legal or certification failure on top of one.

STAGE: MECHANICAL DESIGN

ENCLOSURE DESIGN FOR MASS PRODUCTION

That enclosure you 3D printed is likely impossible to mass produce. Your printer will happily build shapes all day long that no factory could ever produce in volume, and these problems usually cannot be patched at the end. I am an electronics engineer, but when I was building my own product I taught myself enclosure design, and I have seen plenty of designs that had to be completely redesigned before they could be molded. You need to know these rules while you are designing the part, because the last one multiplies your entire manufacturing bill.

IN THIS PART

12.1	Never Leave the Walls Without a Draft Angle	p. 111
12.2	Never Design Non-Uniform Wall Thickness	p. 112
12.3	Never Thin the Walls Without Adding Support Ribs	p. 113
12.4	Never Leave Sharp Internal Corners	p. 114
12.5	Never Ship Snap Fits Tuned on Your Printer	p. 115
12.6	Never Design as if the Mold Leaves No Marks	p. 116
12.7	Never Ignore How the Plastic Flows	p. 117
12.8	Never Accept an Undercut Without Trying the Fixes	p. 118
12.9	Never Design More Plastic Pieces Than You Need	p. 119

12.1 Never Leave the Walls Without a Draft Angle

PART EJECTION

WHAT GOES WRONG IN THE MOLD

- Mass producing a plastic enclosure means injection molding, where melted plastic gets injected into a steel mold, and that one mold can produce hundreds of thousands of parts.
- Your 3D printed enclosure almost certainly has perfectly vertical walls, which was no problem for your printer.
- In a mold, every part has to slide out, and a vertical wall drags against the steel the whole way out, scratching the surface and sometimes refusing to eject at all.
- The molder will add draft for you if you do not, and the walls will end up where they decide, not where your PCB and your buttons need them to be.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Add a degree or two of draft to every wall	Draft is just a slight taper on every wall that runs in the direction the mold opens. One to two degrees is the usual starting point.
Add more draft on textured surfaces	If the surface is textured instead of smooth and polished, you will need a little more draft for the part to release cleanly.
Apply it in CAD from the start	Set the pull direction in your CAD model and use its draft analysis tool, so the draft is designed in rather than added by the molder.

BEFORE YOU COMMIT

- **Draft analysis run:** Your CAD tool's draft analysis shows no undrafted faces in the pull direction, inside and outside.
- **Texture allowance:** Every textured surface has the extra draft the texture depth requires, which your molder can specify.
- **Internal features included:** Ribs, bosses, and pockets have draft too, not just the outer walls.

One mold, hundreds of thousands of parts

Every rule in this part exists because a mold has to release the part cleanly every single cycle, and a problem like this usually cannot be patched at the end.

12.2 Never Design Non-Uniform Wall Thickness

SHRINKAGE

WHAT GOES WRONG IN THE MOLD

- Injection molded plastic shrinks as it cools, and thick sections cool a lot slower than thin ones.
- That uneven cooling is what makes parts warp, and it creates sink marks, those small dents that show up on the surface behind every thick feature. Really thick sections can even form voids, hollow pockets inside the plastic where the material pulled apart as it shrank.
- Your printer never showed you any of this, because a printed part never goes through that kind of packing and shrinking. A chunky 3D printed design comes out perfect, and the exact same design comes out of a mold warped and dented.
- Thick sections also take longer to cool in the mold, which means a longer cycle time and a higher price on every part for the life of the product.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Keep the main wall thickness uniform	Usually in the 2 to 3 mm range for most plastics, and consistent everywhere on the part.
Core out thick areas from the back	If you need a thick area for strength, do not make it one solid mass. Core it out so the walls stay consistent.
Get the strength back with ribs	Ribs give you the stiffness without the thick section, as entry 12.3 covers.

BEFORE YOU COMMIT

- **Thickness analysis run:** Your CAD tool's wall thickness analysis shows the whole part within your target range, with no hot spots.
- **Bosses and corners checked:** Screw bosses and corners, the usual thick spots, have been cored or radiused so they match the wall.
- **Material confirmed:** The wall thickness range matches the plastic you will mold, since the safe range differs between ABS, polycarbonate, and nylon.

The dent behind every thick feature

A sink mark always lands on the cosmetic side, directly opposite the thick feature that caused it, which is why bosses and ribs show through the face of a product.

12.3 Never Thin the Walls Without Adding Support Ribs

STIFFNESS

WHAT GOES WRONG IN THE MOLD

- Your 3D printed prototype was probably built with thick walls and solid sections, because thick sections cost you nothing but a little extra filament.
- Then you thin everything down to uniform 2 to 3 mm walls for molding, and the whole part turns into a thin hollow shell. On top of that, molded ABS is noticeably less stiff than the PLA most people print with.
- So a design that feels rock solid coming off your printer can flex like a cheap toy coming out of a mold.
- Make a rib too thick and it creates the exact sink mark problem from entry 12.2, right on your cosmetic surface, because the spot where a rib meets the wall is thicker than either one by itself.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Add ribs on the inside of the part	Thin walls of plastic that act like structural beams, giving you the stiffness back without thickening the walls.
Add gussets at tall features and corners	Small triangular braces do the same job for tall bosses and corners that need support.
Keep each rib to about half the wall thickness	Ribs are the one place you want thinner plastic on purpose, so a thin rib keeps the junction from turning into a thick section.

BEFORE YOU COMMIT

- **Rib thickness rule applied:** Every rib is roughly half the thickness of the wall it attaches to, with a small radius at the base.
- **Rib height and draft:** Rib height stays within about three times the wall thickness, and each rib has draft on both faces.
- **Stiffness tested in the real plastic:** A CNC-machined or molded sample in the production material has been flexed against a defined deflection limit, since a printed part tells you nothing about molded stiffness.

Thinner on purpose

It sounds like it breaks the uniform wall rule, but ribs are the one place you want thinner plastic on purpose, because a thin rib keeps the junction with the wall from becoming a thick section.

12.4 Never Leave Sharp Internal Corners

STRESS AND TOOLING

WHAT GOES WRONG IN THE MOLD

- Sharp corners are a double mistake, one in the plastic part and one in the steel mold.
- In the plastic part, a sharp inside corner concentrates stress, so when the part gets dropped or flexed, that corner is where the crack starts.
- In the steel mold, a sharp corner on your part means a sharp corner cut into the steel, which costs more to machine and wears out faster than a rounded one.
- Your printer renders corners fine either way, so nobody ever thinks about them until the first drop test or the first tooling quote.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Round every corner, inside and out	A good starting rule is an inside radius of at least half your wall thickness.
Match the outside radius to the wall	Make the outside radius equal to the inside radius plus the wall thickness, so the wall stays constant as it wraps around the corner.
Check the corner is not a thick section	A corner with the inside sharp and the outside rounded is a thick section unless the outside radius equals the wall thickness, and a thick corner brings back the sink and warp problems from entry 12.2.

BEFORE YOU COMMIT

- **Fillets everywhere:** A search of the CAD model for sharp inside edges in the pull direction comes back empty.
- **Wall thickness around corners:** The thickness analysis shows the corners at the same thickness as the walls.
- **Drop test in the real material:** A sample in the production plastic has been dropped onto a hard floor from the height a customer would drop it, and inspected at every corner.

Where the crack starts

When a molded part fails a drop test, the crack almost always begins at a sharp inside corner. Rounding it costs nothing and moves the stress somewhere the part can take it.

12.5 Never Ship Snap Fits Tuned on Your Printer

FASTENING

WHAT GOES WRONG IN THE MOLD

- Everyone has done this: print a part, test the snap fit, tweak the model by a fraction of a millimeter, and print it again until it closes with a perfect click. All of that tuning was specifically for your printer and the filament or resin you printed with.
- Production molded ABS or polycarbonate has different stiffness, different friction, and different shrinkage, so the snap that works perfectly on your printer might not close at all in the production plastic.
- Since printed snaps tend to break along the layer lines, people beef them up to survive on the printer, which leaves them too stiff to flex properly once they are molded in tougher plastic.
- A member of my Hardware Academy ran into exactly this issue, and after our review he decided to go with three plain screw holes instead, to reduce the amount of snap fit optimization needed.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Use screws if you do not absolutely need snaps	Snap fits can take a lot of iterations to get right even at the mold level, so if you do not have to have them, screws tend to be simpler.
Test critical snaps in the actual production plastic	CNC machining is a good way to do that, because you can have a prototype machined from real ABS or polycarbonate and get a much better feel for the production snap.
Design the snap steel-safe	Removing tool steel is cheap and adding it back is not, so design the snap so the expected adjustment after first samples means cutting more steel away, not welding it back.

BEFORE YOU COMMIT

- **Snap count justified:** Every snap fit has a reason a screw would not do, like a customer-facing battery door or a no-visible-fastener requirement.
- **Deflection calculated:** Each snap's beam deflection has been checked against the strain limit of the production plastic, not the printed one.
- **Cycle count:** Snaps that the customer will open and close have been rated for that number of cycles in the production material.

The plastic is the variable

The same geometry behaves differently in PLA, resin, ABS, and polycarbonate. A snap is only proven when it has been tested in the material it will be molded from.

12.6 Never Design as if the Mold Leaves No Marks

COSMETICS

WHAT GOES WRONG IN THE MOLD

- Every injection molded part comes out of the mold with marks on it. There is a parting line where the two mold halves meet, small circles where the ejector pins pushed the part out, and a scar where the plastic was injected.
- Those marks are not a defect, they are just part of the process, and they have to land somewhere.
- Your printer leaves none of these, so people design as if they do not exist. Then the first molded samples show up with a parting line running straight across the product's face and ejector circles right where the customer's thumb rests.
- Moving a parting line after the mold is cut is a new mold, so this is decided once.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Run the parting line along a natural edge	Put it where the design already has an edge, a step, or a texture change, so it reads as a feature instead of a flaw.
Keep ejector pins and the gate on hidden surfaces	The inside of the part, the bottom, or a surface the customer never sees.
Settle all of it with your molder before the design is final	Send the model and ask where they would put the parting line, the gate, and the pins, and adjust the design while it is still cheap to.

BEFORE YOU COMMIT

- **Parting line drawn on the model:** The parting line is marked on the CAD model and reviewed against every cosmetic surface.
- **Ejector plan from the molder:** You have the molder's proposed ejector pin locations and have confirmed none land on a visible or touched surface.
- **Gate location agreed:** The gate location is agreed in writing, and entry 12.7 covers what it changes.

Decide, or the molder decides for you

Every one of these marks will land somewhere. The only question is whether you chose the spot while the design was still in CAD, or found out where it went from the first molded sample.

12.7 Never Ignore How the Plastic Flows

FILL AND WELD LINES

WHAT GOES WRONG IN THE MOLD

- Molten plastic enters the mold through a small opening called the gate, and from that point it has to race outward and fill every corner of the cavity before it freezes.
- Put the gate far away from a thin section, and the plastic solidifies before it gets there, leaving you with what molders call a short shot, a part that is literally incomplete.
- When the flow splits around a hole and meets up again on the other side, it forms a weld line, a visible seam that is also the weakest spot in the part.
- Weld lines behind a screw boss or a button hole are where a part breaks under load, so the flow decides strength as well as looks.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Think about the gate while you are still designing	Decide where the gate can go, and keep the path from the gate to the farthest thin feature as short as you can.
Steer the weld lines	Moving the gate also moves where the weld lines land, so a conversation with your molder before the design is final can keep them off the surfaces your customer sees.
Ask for a flow simulation on a complex part	Many molders will run a mold flow analysis for a modest fee, which shows short shots and weld lines before any steel is cut.

BEFORE YOU COMMIT

- **Flow length ratio:** The distance from the gate to the farthest point, divided by the wall thickness, is inside the range your molder gives for the material.
- **Weld lines mapped:** Every hole and every boss has been checked for the weld line it creates and where it lands.
- **Gate scar hidden:** The gate location from entry 12.6 leaves its scar on a hidden surface.

A seam is a weak spot

A weld line is not just cosmetic. It is where two fronts of cooling plastic met without fully fusing, and it is the first place a stressed part lets go.

12.8 Never Accept an Undercut Without Trying the Fixes

TOOLING COST

WHAT GOES WRONG IN THE MOLD

- An undercut is any feature that traps the part in the mold, like a snap fit, a side hole, or a groove around the outside of the part. Your printer builds these without any problems, because 3D printing does not require the part to be removed from a mold.
- A basic mold only opens in one direction, and an undercut physically locks the part inside.
- Molders can solve this with side actions, extra pieces of the mold that slide in from the side and pull out of the way before the part ejects. They work, but every side action adds cost to the mold, adds a mechanism that wears out, and adds one more thing that can go wrong on every single cycle.
- Side actions also constrain where the parting line and the ejectors can go, so one undercut ripples through every other decision in this part.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Redesign the feature so it is not an undercut	This is possible more often than people think. Move the feature to a surface in the pull direction, or split it across the parting line.
Form it through a hole in the wall below	Put a hole in the wall directly below the undercut, and the mold can form the feature through that hole with no side action needed.
Accept a side action only for what is left	If an undercut survives both workarounds, price the side action and its maintenance before you agree to it.

BEFORE YOU COMMIT

- **Undercut analysis run:** Your CAD tool's undercut detection has been run in the pull direction, and every hit has a plan.
- **Side action count in the quote:** The mold quote lists every side action separately, so you can see what each undercut costs.
- **Commercial products studied:** You have looked under the snap fits on a few commercial products and seen the holes, so you know the trick works.

The hole under the snap

That one trick has saved people thousands of dollars in tooling, and once you know about it, you will start spotting those holes under the snap fits in commercial products everywhere.

12.9 Never Design More Plastic Pieces Than You Need

TOOLING MULTIPLIER

WHAT GOES WRONG IN THE MOLD

- Every mistake so far relates to a single part. This one multiplies across your whole product, and it does the most financial damage.
- Your printer makes a five-piece enclosure feel free, because printing five parts costs about the same as printing one. But every separate plastic piece in your design needs its own mold, and every mold costs thousands of dollars, often tens of thousands.
- The molds for my own product cost over 100,000 dollars. My manufacturer believed in the product enough to finance that cost themselves and amortize it by charging me an extra dollar per unit for the first 100,000 units. That deal saved my launch, but it taught me exactly what every extra plastic piece costs.
- Each extra piece also adds an assembly step, a tolerance stack, and a seam the customer can see.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Design for the fewest separate plastic pieces possible	Combine parts wherever you can, and question every piece that exists only because the printer made it easy.
Use a stock part wherever one will do the job	Off-the-shelf buttons, light pipes, feet, and battery doors mean molds you never have to pay for.
Ask about a family mold for small pieces	If you are stuck with a few small pieces, a family mold forms several different parts in one mold, as long as they are all about the same size.

BEFORE YOU COMMIT

- **Piece count with a reason for each:** Every separate molded piece has a written reason it cannot be combined with another.
- **Mold budget per piece:** You have a tooling quote for each piece, so the cost of every extra part is a number rather than a feeling.
- **Stock alternatives searched:** Catalog parts have been checked for every small piece before it was designed as custom plastic.

The multiplier

A rule that costs you a little on one part costs you five times as much on five parts. This is the entry that decides whether the tooling bill is thousands or tens of thousands.

STAGE: SOURCING AND PRODUCTION

SOURCING AND COUNTERFEIT PARTS

Online marketplaces are a great resource for a wide variety of components, but you have to be careful about the quality of what you actually get. The worst part is that a counterfeit component almost never fails on the day you solder it down. It fails once your product is in a customer's hands, and at that point you are paying for a whole production run over a two dollar part. A fake does not announce itself, and you are already debugging your own design, so the last thing you need is to be debugging your parts at the same time. This part ends with the red flags for spotting a fake listing and the five-minute tests to run on parts before you trust them.

IN THIS PART

13.1	FTDI USB-to-Serial Chips from Unauthorized Sellers	p. 121
13.2	Sensor Modules Sold as Something They Are Not	p. 122
13.3	Cloned STM32 Microcontrollers	p. 123
13.4	Counterfeit and No-Name Electrolytic Capacitors	p. 124
13.5	Counterfeit Power MOSFETs	p. 125
13.6	Rewrapped High-Capacity 18650 Cells	p. 126
13.7	Counterfeit Fuses and Safety-Rated Capacitors	p. 127
13.8	Two-Dollar Mains USB Chargers and Power Supplies	p. 128
13.9	Counterfeit and Clone ICs in General	p. 129
13.10	Single-Source and End-of-Life Components	p. 130

SECTION

PART 13

DOC NO.

NVR-01

ENTRIES

10

SOURCE

PREDICTABLEDESIGNS.COM

13.1 FTDI USB-to-Serial Chips from Unauthorized Sellers

COUNTERFEIT ICS

WHY IT FAILS IN A PRODUCT

- The FT232 from FTDI is probably the most counterfeited chip in hobby electronics.
- Back in 2014, FTDI pushed out a Windows driver update that detected the clones and reprogrammed their USB product ID to zero, which left them dead as far as Windows was concerned. A later version stopped killing them and instead made them send the text "non genuine device found" out the serial port, which corrupted the data going to whatever was listening.
- Now picture that being your product on a customer's desk, and it stops talking to their computer after a routine Windows update they did not even choose to install. You changed nothing, but you are the one taking the support call and paying the return shipping.
- The cables are worse than the chips, because a cable is a sealed assembly and you cannot even read the marking on the part inside it.

BUY THIS INSTEAD

THE SOURCE OR REPLACEMENT	WHY THIS WORKS
Genuine FTDI from an authorized distributor	Buy the FT232 from Digi-Key, Mouser, or another authorized distributor, so you know you are getting an authorized FTDI part.
WCH CH340 through legitimate channels	If cost is what is pushing you toward the cheap option, the CH340 is a perfectly good part when you buy it through a legitimate channel, with drivers built into current operating systems.
Silicon Labs CP2102 or CP2104	Another solid bridge from a major vendor, well supported and stocked by every authorized distributor.
Native USB on the microcontroller	The ESP32-S3 and most STM32 parts have USB built in, which removes the bridge chip and the counterfeit risk entirely.

BEFORE YOU COMMIT

- **Distributor invoice on file:** The bridge chip on your bill of materials has an authorized distributor as its only approved source, and the invoice proves it.
- **Driver test on a clean machine:** Plug a production unit into a freshly updated Windows and macOS machine and confirm it enumerates with the vendor's current driver.
- **Cables specified too:** If the product ships with a USB-to-serial cable, it is sourced by part number from FTDI or an authorized distributor, not from a marketplace listing.

The update you did not install

A counterfeit bridge chip turns every operating system update into a potential support call, for the whole life of the product, over a part that saved a dollar.

13.2 Sensor Modules Sold as Something They Are Not

MISLABELED PARTS

WHY IT FAILS IN A PRODUCT

- The BME280 breakout board is the poster child for this one. The BME280 from Bosch measures temperature, pressure, and humidity, while the cheaper BMP280 measures temperature and pressure with no humidity sensor at all.
- Plenty of listings sell BMP280 modules with BME280 printed in the title, and the two boards look identical.
- You wire it up, run a standard BME280 library, and it refuses to start, because that library checks the chip's ID and gets the wrong number back. Then you go hunting through your wiring, your pull-up resistors, and your I2C address, when the real problem is that the chip is not the one you ordered.
- The same trick shows up across sensor families, with older or lower-grade parts sold under the newer part number, and the module form factor hides it.

BUY THIS INSTEAD

THE SOURCE OR REPLACEMENT	WHY THIS WORKS
The sensor through an authorized channel	For anything going into a product, buy the module, or better the bare sensor, through an authorized distributor.
The sensor on your own board	Put the sensor on your own PCB where you control the sourcing, the placement, and the layout around it, as Part 4 covers.
The ID register as your acceptance test	That same ID register tells you exactly what you are holding. A real BME280 comes back as hex 60, while the humidity-free BMP280 reads hex 58.

BEFORE YOU COMMIT

- **ID read on every reel:** The first parts from every new reel or batch get their ID register read on the bench before assembly starts.
- **Library check kept in firmware:** The production firmware checks the sensor ID at boot and logs a fault if it is wrong, so a substitution at the assembler is caught on the line.
- **Part number, not description:** The bill of materials carries the manufacturer part number, and the assembler is not allowed to source by description.

Read the register first

Before you touch the wiring or the pull-ups, read the chip ID. It takes thirty seconds and it catches the obvious substitutions, though a matching ID does not prove the part is genuine.

13.3 Cloned STM32 Microcontrollers

COUNTERFEIT ICS

WHY IT FAILS IN A PRODUCT

- The STM32F103 is one of the most cloned microcontrollers on the planet, and what shows up on a lot of those cheap boards did not come from ST at all. Sometimes it is a pin-compatible part from another manufacturer with ST's markings reprinted on top, and sometimes it is a lower-grade part remarked as a higher-grade one.
- Those chips almost always boot, they run your blinking LED demo just fine, and then they drift from the datasheet in small ways: timer behavior that is a little off, a peripheral that does not do what the reference manual says, or a flash size that does not match what the part number claims.
- You end up tearing apart your firmware for weeks when your firmware was never the problem.
- There is a second issue that matters even more for a product, since you cannot reorder these with any confidence when the next batch could easily be a different die.

BUY THIS INSTEAD

THE SOURCE OR REPLACEMENT	WHY THIS WORKS
Digi-Key, Mouser, or Arrow	Buy your microcontrollers from an authorized distributor, which covers you from prototypes all the way through your early production runs.
A volume quote from your distributor	Once your volumes climb into the tens of thousands, ask your distributor for a volume quote, because nobody in production keeps paying catalog price.
ST directly at scale	At real volume, talk to ST directly. The pricing is better than any marketplace and the parts are the parts.

BEFORE YOU COMMIT

- **Device ID and flash size read:** Read the device ID register and the flash size register on a sample from each batch and compare them with the part number you ordered.
- **Marking inspected:** Under magnification, the top of the package shows a clean laser mark with no sign of sanding or overprinting, and the date codes match across the reel.
- **Single approved source:** The bill of materials lists the authorized distributor as the only approved source for the microcontroller.

Debugging the wrong thing

A clone that boots is worse than one that does not, because it sends you into weeks of firmware debugging for a problem that lives in the silicon.

13.4 Counterfeit and No-Name Electrolytic Capacitors

COUNTERFEIT PASSIVES

WHY IT FAILS IN A PRODUCT

- When you see Panasonic, Vishay, or other premium-brand capacitors selling for a fraction of the distributor price, what you are actually buying is a plastic shell with a respected brand name printed on it. That shell is the easiest thing in the world to fake, and what is underneath is usually a low-grade cap with the wrong ESR and a much shorter life.
- ESR, the internal resistance of the capacitor, controls how much ripple current it can handle before it heats up and dries out. Get that wrong in a switching power supply and your product does not fail on your bench, it fails eight months later across your production run.
- The capacitor plague of the early 2000s, when bad electrolyte took out motherboards by the millions, is not ancient history. It is still happening with bottom-tier capacitor brands, which dry out early, drift out of spec, and kill a product's reliability over one to three years.
- If your contract manufacturer is substituting mystery-brand caps into your bill of materials without your approval, that is a BOM control problem you need to fix immediately.

BUY THIS INSTEAD

THE SOURCE OR REPLACEMENT	WHY THIS WORKS
Panasonic, Nichicon, or Rubycon from authorized distributors	This is one component category where brand matters, and the authorized channel is the only way to know the brand is real.
A quality cap from a budget brand over a fake premium one	If your budget is tight, a real capacitor from a lesser brand will always beat a counterfeit version of a premium brand.
Ceramic or polymer where the design allows	Switch to ceramic or polymer aluminum capacitors where you can, as entry 7.5 covers. The dry-out problem goes away with the liquid, though the sourcing rules still apply.

BEFORE YOU COMMIT

- **Approved vendor list at the assembler:** Your assembler has a written approved manufacturer list for every capacitor line, and substitutions require your sign-off.
- **ESR measured on samples:** An ESR meter reading on a few parts from each batch matches the datasheet within its tolerance.
- **Price sanity check:** Any capacitor priced well below the authorized distributor's volume price is treated as counterfeit until proven otherwise.

Eight months later

A capacitor with the wrong ESR passes every test on the bench. It fails after the warranty period has started and before it has ended, across the whole run at once.

13.5 Counterfeit Power MOSFETs

COUNTERFEIT POWER PARTS

WHY IT FAILS IN A PRODUCT

- A MOSFET is just an electronic switch, and how much current it can safely carry is one of its most important specs. Counterfeit power MOSFETs are typically a smaller, cheaper die inside the right-looking package with the right part number etched on top.
- It will pass a bench test at low current every time, and it may even handle your full rated current at room temperature. Then you run it inside a sealed enclosure on a hot day, that undersized die cannot shed the heat, and it fails.
- Even when it survives, running a small die past what it was built for damages it a little at a time, so the part dies months later in the field instead of on your bench.
- This one bites hardest in motor drivers, battery chargers, and anything else switching real power, which is exactly where a failure becomes smoke.

BUY THIS INSTEAD

THE SOURCE OR REPLACEMENT	WHY THIS WORKS
Authorized distributors, full stop	Power parts are the last place to save fifty cents, so buy them from authorized distributors with no exceptions.
A load test on any inherited design	If you inherit a design with mystery FETs already in it, load test one at its rated current under worst-case conditions before you trust it across a production run.
A part with margin	Choose a MOSFET rated well above your peak current and voltage, so even a marginal part has room, and the thermal design in entry 10.4 has margin too.

BEFORE YOU COMMIT

- **On-resistance measured:** Measure the on-resistance of samples at a known gate voltage and compare it with the datasheet. A small die reads high.
- **Thermal test at rated current:** Run the part at full load in the real enclosure at the highest ambient temperature and watch the case temperature with a thermal camera.
- **Source locked:** The bill of materials names an authorized distributor as the only source for every power semiconductor.

The last place to save fifty cents

A counterfeit power part does not fail politely. It fails hot, in the customer's hands, and in the part of the circuit that carries the most energy.

13.6 Rewrapped High-Capacity 18650 Cells

COUNTERFEIT BATTERIES

WHY IT FAILS IN A PRODUCT

- You will see listings for 9,900 mAh 18650 cells, sometimes even 12,000 mAh, for a couple of bucks each. Physics does not allow that, since the best cells Samsung, LG, and Panasonic can build in that can size top out around 3,500 mAh.
- Anything claiming way more than that is a rewrapped cell, often a used pull from a scrapped laptop pack, and the measured capacity is frequently under 1,000 mAh.
- The fake capacity is the smaller problem, because you also have no idea what that cell has been through, what its internal resistance looks like now, or whether it has any protection circuitry at all.
- Putting a cell like that inside a product you are selling means you are knowingly carrying a fire risk, and the recall and liability that follow are yours alone.

BUY THIS INSTEAD

THE SOURCE OR REPLACEMENT	WHY THIS WORKS
Vetted distributors with traceability	Buy your cells from distributors who can trace them back to the actual manufacturer, with the cell datasheet and the UN38.3 summary to match.
Top-tier cell makers only	Samsung, LG, Panasonic, Molicel, and EVE publish real datasheets, and a cell you cannot trace to one of them does not belong in a product.
Capacity testing as an incoming check	Capacity test every cell, or a statistical sample of every batch, on a discharge tester before it goes into a build.

BEFORE YOU COMMIT

- **Physics check:** Any 18650 listing well above 3,500 mAh that cannot be traced to a real manufacturer's datasheet is treated as fake.
- **Weight and wrap:** Genuine high-capacity 18650s weigh in the mid-40-gram range, a protected cell runs a few millimeters longer than a bare one, and the wrap carries the maker's printed model number, not a generic brand.
- **Internal resistance logged:** Internal resistance is measured on incoming cells and compared with the datasheet, since a tired cell reads high long before it fails.

See Part 5

Fake cells are just one of the ways a battery can sink your product. Part 5 covers the rest, from protection to the documentation your freight forwarder will ask for.

13.7 Counterfeit Fuses and Safety-Rated Capacitors

SAFETY PARTS

WHY IT FAILS IN A PRODUCT

- Fuses and safety capacitors have one job, which is to fail correctly when something else goes wrong. A fuse is supposed to open at a specific current within a specific time, so it cuts power before anything else overheats.
- Safety-rated X and Y capacitors sit right at your mains input, which puts them in the most dangerous spot on the board. An X capacitor goes across the two mains lines, a Y capacitor connects a mains line to ground or to something your customer can touch, and both are built to fail open rather than short.
- Counterfeit versions of both are everywhere, and the safety agency marks printed on them are just ink. A fake fuse that holds at three times its rating is a disaster waiting to happen, and a fake Y capacitor that shorts instead of opening puts mains voltage somewhere a person can touch.
- There is no bench test that rescues you either, since a safety part only proves itself during the one event you hope never happens.

BUY THIS INSTEAD

THE SOURCE OR REPLACEMENT	WHY THIS WORKS
Authorized distributors for every safety part	This is the one category where I would not even weigh the trade-off. Buy fuses and safety capacitors from authorized distributors only.
Verify the marks are real	Look up the agency file number for the exact part on the certification body's database, and confirm the manufacturer and the rating match.
Keep the paperwork	The certificates, the distributor invoices, and the lot numbers go in the compliance file, because the test lab and any future investigation will ask for them.

BEFORE YOU COMMIT

- **File number looked up:** Every fuse and every X or Y capacitor has its agency file number verified on the certification database, not just printed on the part.
- **Safety-critical flag on the BOM:** Safety parts are marked no substitution on the bill of materials, with the authorized source named.
- **Traceability per lot:** You can trace the lot number of the safety parts in any given production run back to the distributor invoice.

One event, no second chance

A safety part is the only component you buy hoping never to find out whether it works. That is why the paperwork is the test.

13.8 Two-Dollar Mains USB Chargers and Power Supplies

MAINS POWER

WHY IT FAILS IN A PRODUCT

- These are the two dollar wall adapters that plug straight into an outlet, and teardowns of the cheap ones turn up the same problems over and over.
- The isolation barrier between the mains side and the low-voltage side measures under 1 mm in some of them, when safety standards call for several millimeters of clearance.
- The certification marks molded into the case were never earned by anybody, and basic input protection like a fuse or a surge suppressor gets left off to save a few cents.
- Buying one for your own bench is your call, but designing your product around one is where I would draw a hard line. If that supply fails, dangerous AC voltage may make it to the low-voltage side, and that puts your customer at real risk.

BUY THIS INSTEAD

THE SOURCE OR REPLACEMENT	WHY THIS WORKS
A certified supply from a real brand	Use a supply from a manufacturer that will hand you an actual test report with a certificate number you can look up yourself.
The certificate number verified	Search the certificate number on the certification body's database and confirm it covers the exact model, the input range, and the output rating you are shipping.
Regional variants specified	Plug types, input voltage ranges, and efficiency rules differ by market, so specify the certified variant for each country you sell in.

BEFORE YOU COMMIT

- **Certificate in the compliance file:** The supply's test report and certificate are on file and match the model number printed on the adapter you ship.
- **Teardown of one sample:** Open one adapter from the first batch and measure the isolation spacing and confirm the fuse is present, so the certificate describes the unit in your box.
- **Supply in the product certification:** The adapter is listed in your own product's technical file, since the product is certified as a system that includes it.

The hard line

A cheap adapter on your own bench risks your bench. A cheap adapter in your product risks your customer, and that is the one risk you cannot buy back.

13.9 Counterfeit and Clone ICs in General

SUPPLY CHAIN

WHY IT FAILS IN A PRODUCT

- Fake STM32s, cloned FTDI chips, and suspiciously cheap "original" components from unauthorized distributors are more common than most people realize.
- They pass bench testing just often enough to make it into a production run, and then they fail in the field at scale when it is far too late and far too expensive to fix.
- Common signs include pricing that seems too good to be true, packaging inconsistencies, and unusual behavior under edge conditions like temperature extremes or the corners of the supply voltage range.
- If you are sourcing ICs from unauthorized marketplaces to save a few cents per chip, you are gambling your entire production run, and during a shortage the temptation to gamble is strongest exactly when the counterfeit supply is largest.

BUY THIS INSTEAD

THE SOURCE OR REPLACEMENT	WHY THIS WORKS
Authorized distributors for every IC	Mouser, Digi-Key, Arrow, and Farnell buy directly from the manufacturer, so the part in the reel is the part on the label.
The manufacturer's authorized distributor list	Every chip vendor publishes its authorized distributors. If a seller is not on that list, they are not authorized, whatever their listing says.
A broker only with testing, during a shortage	If a shortage forces you to a broker, use one that offers independent component testing and full traceability, and test the parts before assembly.

BEFORE YOU COMMIT

- **Every IC has an authorized source:** The bill of materials lists an authorized distributor for every integrated circuit, and the assembler buys from that list.
- **Incoming inspection on new lots:** The first parts from any new lot get the checks at the end of this part before they go on the line.
- **Red flags reviewed:** Every purchase outside the authorized channel has been screened against the red flags at the end of this part, in writing.

Just often enough

The dangerous counterfeit is not the one that fails on the bench. It is the one that passes just often enough to reach production and fails at scale in the field.

13.10 Single-Source and End-of-Life Components

LIFECYCLE

WHY IT FAILS IN A PRODUCT

- If only one factory on earth makes your critical component, you are one supply disruption away from a dead product. The chip shortage of a few years ago made this painfully obvious, and it will happen again.
- End-of-life components are even worse, because you are designing on a countdown timer that you might not know about.
- This is especially dangerous for startups with long development timelines, because your key part can go end-of-life before you even ship your first unit, as the MPU-9250 story in entry 4.5 shows.
- If no second source exists, that is a risk you need to consciously accept up front, not discover later when you are trying to place a production order.

DO THIS INSTEAD

THE FIX	HOW TO APPLY IT
Check lifecycle status before committing any part	Every manufacturer publishes a lifecycle status for each part, and the distributors flag it too. Read it before the part goes on the bill of materials.
Look for second-source alternatives	For passives, connectors, and standard logic, specify a second manufacturer on the bill of materials. For ICs, note the nearest pin-compatible alternative.
Document the single-source risks you accept	For the parts with no alternative, write down the risk, the last-time-buy plan, and the redesign cost, so it is a decision rather than a surprise.

BEFORE YOU COMMIT

- **Lifecycle column on the BOM:** The bill of materials has a lifecycle status column, and every part is active or has a documented plan.
- **Change notifications subscribed:** You receive product change and end-of-life notifications from the distributor or the manufacturer for every IC in the design.
- **Stock strategy:** For any single-source part, you know the lead time and hold enough stock to cover it, or you have budgeted a last-time buy.

A countdown you did not know about

A part marked not recommended for new designs is living on borrowed time, and no amount of tutorial popularity will bring it back.

Red Flags and Incoming Inspection

Spot a fake listing in ten seconds, and test the parts in five minutes before you trust them.

Every entry in this part comes down to two habits. Screen the listing before you buy, and test the parts before they go on the line. Together they catch the obvious counterfeits, though neither one proves a lot is genuine. That is what the authorized channel is for.

RED FLAGS: SPOT A FAKE LISTING IN TEN SECONDS

- ☐ **Price far below the authorized channel:** A part selling for a fraction of the distributor's volume price is counterfeit until proven otherwise, because nobody sells genuine parts at a loss.
- ☐ **A premium brand at a commodity price:** Panasonic capacitors, Samsung cells, or ST microcontrollers at no-name prices mean the brand is the only genuine thing in the listing.
- ☐ **Specs physics does not allow:** A 9,900 mAh 18650, a fuse with no agency file number, or a MOSFET rated above its own datasheet.
- ☐ **Certification claims with no paperwork:** FCC, CE, or UL in the title with no FCC ID, no certificate number, and no test report on request.
- ☐ **No manufacturer, no datasheet, stock photos:** A seller who cannot name the manufacturer or send the real datasheet is not selling the real part.
- ☐ **The word original doing all the work:** Authorized distributors do not need to say original. Sellers who say it in every title are answering a question you should keep asking.
- ☐ **Packaging that does not match the vendor:** Loose parts in bags instead of reels or trays, missing moisture barrier bags on moisture-sensitive parts, and labels without lot or date codes.

FIVE-MINUTE INCOMING INSPECTION TESTS

- ☐ **Read the ID register:** Sensors, microcontrollers, and most digital ICs report a device ID. Read it on the first parts from every batch and compare it with the datasheet.
- ☐ **Inspect the marking under magnification:** Look for sanded or blacktopped package tops, uneven laser marking, and date codes that do not match across the reel.
- ☐ **Measure the ESR:** An ESR meter reading on a few electrolytics from each batch, compared with the datasheet, catches the wrong electrolyte before it goes on the board.
- ☐ **Measure on-resistance and heat:** For power MOSFETs, measure the on-resistance at a known gate voltage and run one part at your real load while watching its temperature.
- ☐ **Weigh the cells and start a capacity test:** A genuine high-capacity 18650 weighs in the mid-40-gram range. Start a full discharge on a sample, five minutes to set up and hours to run, and compare it with the datasheet.
- ☐ **Look up the certification file number:** For fuses, safety capacitors, wireless modules, and power supplies, verify the file number or FCC ID on the certification body's database, since paperwork is the only test a safety part gets.

Make it a habit, not a project

Run the red flags on every purchase outside the authorized channel and the inspection tests on the first parts from every new lot. Ten minutes per batch is the cheapest insurance in this handbook.

The Master Swap Table

Every entry in the handbook on a few pages, with the page to turn to for the full explanation.

Use this table as the quick reference once you have read the entries. Each row is a link to the full entry, and the page number is the merged page in this PDF.

#	NEVER	INSTEAD	PAGE
Part 1: Microcontrollers and Processors			
1.1	ESP8266 in a new design	ESP32-C3 (or C6, S3 for more capability)	5
1.2	PIC16 for a new product	STM32C0 or TI MSPM0 (Cortex-M0+)	6
1.3	Ultra-cheap no-name MCUs with no documentation	STM32C0 or MSPM0 (CH32V003 as the one exception)	7
1.4	STM8 in a new design	STM32C0 or STM32G0	8
1.5	MSP430 for a new general-purpose design	TI MSPM0 or STM32U0 (MSP430 FRAM parts remain valid)	9
1.6	ATmega328P carried from an Arduino prototype into production	ESP32-C3, STM32C0, or RP2040	10
1.7	A top-tier MCU plus external RAM, flash, and display controller	STM32MP1, Allwinner T113, or a Raspberry Pi Compute Module	11
1.8	Flagship phone-class application processors	STM32 or ESP32, or a Raspberry Pi Compute Module for more horsepower	12
1.9	A custom chip to make the product smaller	Off-the-shelf chips, system-in-package parts, and proven modules	13

#	NEVER	INSTEAD	PAGE
Part 2: Arduino and Development Boards			
2.1	Arduino Uno R3, Mega 2560, or Nano as the product foundation	ESP32 dev board, nRF52840 DK, STM32 Nucleo, or Raspberry Pi Pico	15
2.2	Arduino Nano RP2040 Connect	A single ESP32 dev board, or a Raspberry Pi Pico if you need no wireless	16
2.3	Arduino MKR boards (Wi-Fi, LoRa, cellular)	ESP32 dev board; Quectel or SIMCom cellular; SX126x-based LoRa	17
2.4	Arduino Uno R4 Wi-Fi	An Espressif ESP32-S3 dev board, then a pre-certified ESP32-S3 module	18
2.5	Arduino Portenta H7	STM32H7 Nucleo or Discovery board with STM32CubeIDE	19
2.6	Arduino Uno Q (Qualcomm Dragonwing)	STM32MP1, Allwinner T113, or Raspberry Pi Compute Module 5	20

#	NEVER	INSTEAD	PAGE
Part 3: Wireless Modules and Connectivity			
3.1	HC-05 and HC-06 Bluetooth Classic modules	A BLE module on Nordic or Silicon Labs silicon (Fanstel, Raytac, Laird)	22
3.2	HM-10 BLE module	A current BLE 5.x module on Nordic nRF52 silicon	23
3.3	Legacy BLE 4.0/4.1 modules like the RN4020, BLE112/BLE113, and CC2540-based parts	A current BLE 5.x module on modern silicon from a serious vendor	24
3.4	ESP-01 module	A shielded ESP32 WROOM or MINI module (S3, C6, or C3) with modular approval	25
3.5	Generic nRF24L01 clone modules	Decide the protocol first; plain BLE or BLE mesh for most point-to-point and small-mesh products	26
3.6	RF-trace-pin modules and U.FL/SMA modules used with an antenna the grant does not cover	Modules with an integrated antenna and valid modular FCC approval	27
3.7	Generic uncertified LoRa modules	Pre-certified LoRa modules in the right regional variant from Murata or RAK Wireless	28
3.8	SIM800 and SIM900 2G modules	LTE-M, NB-IoT, or LTE Cat-1 bis modules from Quectel, u-blox, or Telit	29
3.9	Listings claiming FCC or CE certification with paperwork that does not hold up	Modules from manufacturers who publish real certification documentation you verify yourself	30
3.10	Designing your own radio, antenna, and matching network	Certified modules and wireless technologies with proven reference designs	31
3.11	5G modules in a product that does not need the throughput	Pre-certified LTE Cat-M1, NB-IoT, or 4G modules	32

#	NEVER	INSTEAD	PAGE
Part 4: Sensors			
4.1	DHT11 and DHT22	Sensirion SHT40/SHT41, TI HDC3020, or Bosch BME280	34
4.2	HC-SR04 ultrasonic module	ST VL53 time-of-flight family, MaxBotix sensors, or a reflective IR sensor	35
4.3	HC-SR501 PIR motion module	Panasonic EKMC or EKMB series, or ST STHS34PF80 for presence	36
4.4	MQ-series gas sensors	Sensirion SGP40/SGP41 or SCD40/SCD41, electrochemical cells, or certified detection modules	37
4.5	MPU-6050 IMU	TDK ICM-42688-P, Bosch BMI270, or ST LSM6 family	38
4.6	Cadmium sulfide photoresistors	Vishay VEML7700, TI OPT3001, or Lite-On LTR-303	39
4.7	ACS712 hobby modules on mains voltage	Shunt with TI AMC1302 isolated amplifier, a current transformer, or Allegro ACS37800	40

#	NEVER	INSTEAD	PAGE
Part 5: Batteries and Charging			
5.1	9V alkaline battery	Two AA cells in series, or a single-cell lithium polymer with a charger IC	42
5.2	Alkaline AA and AAA cells in a high-drain product	Low self-discharge NiMH, or lithium iron disulfide (L91 class) for cold and heavy loads	43
5.3	Coin cells for anything beyond a trickle of current	A lithium iron disulfide AA or a small rechargeable, or a coin cell compartment built to Reese's Law	44
5.4	Sealed-in, non-replaceable batteries	A user-replaceable battery behind a screw-secured door, decided at the concept stage	45
5.5	Bare, unprotected 18650 cells	Protected cells or a pack with a battery management system, from Samsung, LG, Molicel, or EVE	46
5.6	Bare RC-style lithium polymer pouch cells	A pouch cell with an integrated protection board, a proper charge IC, and mechanical clearance	47
5.7	Lithium packs with no datasheet, no test reports, and no traceable cells	Packs from established manufacturers with the UN38.3 summary, cell datasheet, and IEC 62133 report	48
5.8	A custom-shaped pack or an unusual chemistry to squeeze in capacity	Off-the-shelf pre-certified packs, with the enclosure designed around them	49
5.9	TP4056-style charging modules from marketplaces	A charger IC with a thermistor input on your own board: Microchip MCP73833 or the TI BQ series	50
5.10	Wireless charging as a default feature	A USB-C port, pogo pins, or a custom cradle	51

#	NEVER	INSTEAD	PAGE
Part 6: Voltage Regulators and Power Conversion			
6.1	7800 series linear regulators	A buck converter for big step-downs, a modern low-dropout LDO for small ones	53
6.2	LM2596 and the one-dollar modules built on it	A modern synchronous buck at 500 kHz or higher, on your own layout	54
6.3	MC34063 for buck, boost, or inverting supplies	A dedicated buck or boost IC built for the job	55
6.4	No-name switching regulator ICs from marketplace modules	Regulators from TI, Analog Devices, ST, and other major vendors, bought through major distributors	56
6.5	AMS1117 carried over from a dev board	A modern LDO with dropout under 250 mV sized for the load; TPS793 or TPS799 for rails under 200 mA	57
6.6	The cheapest switching module you can find dropped into a production design	A shielded, well-laid-out module from a reputable supplier, or your own layout	58
6.7	HT7333 and HT7533 in anything beyond a trivial load	An LDO with thermal protection, proper load regulation, and specs across temperature	59

#	NEVER	INSTEAD	PAGE
6.8	LM317 (or any adjustable regulator) to make a fixed rail	A fixed-output LDO matched to the rail	60
6.9	MCP1700 picked for its 1.6 uA quiescent current without checking the load profile	An LDO that combines low quiescent current with fast transient response and enough current	61
6.10	Unshielded inductors on a switching regulator	Shielded or semi-shielded inductors from Würth, Coilcraft, or TDK	62

#	NEVER	INSTEAD	PAGE
Part 7: Capacitors, Passives, and Discretes			
7.1	Y5V and Z5U ceramics	X7R or X5R for decoupling and bulk, C0G for anything precise	64
7.2	0402 and 0201 packages by default, and high values crammed into them	0603 as the default, with DC bias curves checked at your operating voltage	65
7.3	A ceramic output capacitor on a regulator whose loop depends on ESR	A capacitor in the ESR range the datasheet calls for, or an LDO stable with ceramics	66
7.4	X7R or X5R in timing, filter, oscillator, and precision analog circuits	C0G (NP0) class 1 ceramics	67
7.5	85 C, 2,000-hour electrolytics placed near heat	105 C parts rated 5,000 hours or more, kept away from heat, or polymer and ceramic replacements	68
7.6	A standard film or ceramic capacitor across the mains line or the isolation barrier	An X2 across the line and a Y1 across the barrier, with the approvals documented	69
7.7	Manganese dioxide tantalums run near their rated voltage	Polymer tantalums, or traditional tantalums derated to half the rated voltage or less	70
7.8	Through-hole resistors, capacitors, and ICs by habit	SMD equivalents, with through-hole reserved for connectors and heavy power parts	71
7.9	Mechanical relays by default	A MOSFET for DC loads, a solid-state relay where control isolation is needed	72

#	NEVER	INSTEAD	PAGE
Part 8: Connectors, Interfaces, and Peripherals			
8.1	USB Micro-B on a new product	USB-C	74
8.2	Barrel jacks for power input	USB-C for most products, a locking connector rated for the current above that	75
8.3	Bare 0.1-inch pin headers for board-to-board or wire connections	Keyed, latching connectors from JST, Molex, or Hirose	76
8.4	MIPI, Thunderbolt, Lightning, and other high-speed or licensed interfaces	SPI, I2C, LVDS, or HDMI	77
8.5	USB 3.0 SuperSpeed or USB Power Delivery when the product does not need them	USB 2.0 for data and 5 V USB for power	78

#	NEVER	INSTEAD	PAGE
8.6	Smartphone-class AMOLED and MIPI DSI panels	A standard TFT LCD with an SPI or parallel RGB interface from Newhaven or BuyDisplay.com	79
8.7	Audiophile-grade analog audio paths without analog expertise	Digital audio over USB, I2S, or a Bluetooth module, and a copied reference design for any analog stage	80

#	NEVER	INSTEAD	PAGE
Part 9: Circuit Protection			
9.1	A power input with no reverse polarity protection	A P-channel MOSFET on the high side, or an ideal diode controller at higher current	82
9.2	Exposed connectors and buttons with no ESD diodes	ESD diodes on every externally accessible line, placed at the connector	83
9.3	A power path with no fuse, PTC, or current limit	A resettable PTC for most products, an eFuse for precision, a one-time fuse for safety rails	84
9.4	A power input with no TVS diode and no overvoltage cutoff	A TVS diode for transients plus an eFuse or OVP IC for sustained overvoltage	85
9.5	Relays, motors, and solenoids driven with no flyback path	A flyback diode at every one-direction DC coil, or the driver's own recirculation path	86
9.6	Lithium batteries with no protection circuitry	Protected cells plus a dedicated charger IC with correct termination and temperature limits	87

#	NEVER	INSTEAD	PAGE
Part 10: PCB Layout			
10.1	Decoupling capacitors placed away from the power pin	Smallest cap closest to the pin, short wide traces, ground vias at the pads	89
10.2	A switching regulator laid out without minimizing the hot loop	Hot-loop cap (input for a buck, output for a boost) tight to the part, loop on one layer, datasheet copied	90
10.3	Power routed at signal trace width through single vias	Traces sized from a width calculator, copper pours for high current, multiple vias in parallel	91
10.4	A power pad with no thermal vias and no copper	The datasheet land pattern, the full via array, and as much copper as you can spare	92
10.5	Vias inside component pads without filled and capped vias specified	Vias next to the pad with a short trace, or filled and capped vias budgeted from day one	93
10.6	Signal traces crossing a gap in the ground or power plane beneath them	Every trace checked against the planes under it on every layer, and rerouted around gaps	94
10.7	USB, Ethernet, and HDMI pairs routed without length matching, impedance control, or a solid reference	The differential pair tools in your software and an impedance-controlled stackup from your fab	95
10.8	Clocks, USB, and memory buses routed along the board edge	Fast signals at least 4 to 5 times the trace-to-plane distance from the edge, on inner layers	96

#	NEVER	INSTEAD	PAGE
10.9	Ground planes on different layers with no stitching vias between them	Ground stitching vias across the board, along the edges, and around high-speed and RF routing	97
10.10	Copper, traces, components, or metal hardware inside the antenna keepout	The datasheet keepout copied exactly, on every layer, with metal hardware kept clear	98
10.11	A board with no test points for production test and debugging	Test pads on every power rail, every bus, and every programming and debug signal	99

#	NEVER	INSTEAD	PAGE
Part 11: Firmware			
11.1	Verbose print logging compiled into the release build	Logging behind levels, with verbose output compiled out and every remaining line audited	101
11.2	Delay calls as the timing architecture	A non-blocking design built on timers, a state machine, or a small RTOS	102
11.3	Saving settings or counters to flash on every loop or every second	Write on change only with a delay, wear leveling for logs, or FRAM for heavy write loads	103
11.4	Shipping with the watchdog disabled, or petted from an interrupt	Independent watchdog enabled and reset from the main loop after every critical task checks in	104
11.5	Brown-out detection disabled to stop nuisance resets	Brown-out detection enabled, threshold set above the flash-write minimum, and the power problem fixed	105
11.6	SWD or JTAG left open on shipped units	Readout protection enabled as the last step of production programming, at a level you can still erase	106
11.7	API keys, cloud secrets, or Wi-Fi credentials baked into the firmware	Provisioned network credentials and unique per-device cloud credentials in protected storage	107
11.8	No update path, or an update path that accepts any image	Signed images verified before boot, two flash copies, and a bootloader that confirms the new image boots	108
11.9	A universal default login like admin and password on every unit	A unique password per unit printed on the label, or a forced password set on first boot	109

#	NEVER	INSTEAD	PAGE
Part 12: Enclosure Design			
12.1	Perfectly vertical walls	A degree or two of draft on every wall in the pull direction, more on textured surfaces	111
12.2	Thick solid sections mixed with thin walls	Uniform 2 to 3 mm walls, with thick areas cored out and strengthened with ribs	112
12.3	Thin uniform walls with nothing to stiffen them	Ribs and gussets at about half the wall thickness	113
12.4	Sharp inside corners on the part	Inside radius of at least half the wall thickness, outside radius equal to inside plus wall	114

#	NEVER	INSTEAD	PAGE
12.5	Snap fits iterated to perfection in printed plastic	Screws where you can, and snaps tested in the actual production plastic where you must	115
12.6	Ignoring the parting line, ejector pin marks, and gate scar	Decide where every mark goes before the design is final, and settle it with the molder	116
12.7	Designing without thinking about the gate and the flow path	A short gate-to-thin-feature path, with weld lines kept off visible surfaces and load paths	117
12.8	Undercuts that force side actions into the mold	Redesign the feature, or form it through a hole in the wall below it	118
12.9	A five-piece enclosure because printing five parts felt free	The fewest separate pieces possible, stock parts where they fit, and a family mold for small pieces	119

#	NEVER	INSTEAD	PAGE
Part 13: Sourcing and Counterfeit Parts			
13.1	FT232 chips and cables from marketplace sellers	Genuine FTDI from Digi-Key or Mouser, or CH340 and CP2102 through legitimate channels	121
13.2	BME280 and similar sensor modules from marketplace listings	The sensor from an authorized channel, on your own board, ID register read as a screen	122
13.3	STM32 parts and boards from marketplace sellers	Authorized distributors for every stage, and a volume quote or ST directly at scale	123
13.4	Premium-brand electrolytics at a fraction of distributor price, or no-name brands from the assembler	Panasonic, Nichicon, or Rubycon through authorized distributors, with BOM control at the assembler	124
13.5	Power MOSFETs from unauthorized sellers	Authorized distributors only, with a load test on anything you inherit	125
13.6	9,900 mAh and 12,000 mAh 18650 listings	Traceable Samsung, LG, Panasonic, or Molicel cells from vetted distributors, capacity tested	126
13.7	Fuses and X or Y capacitors from unauthorized sellers	Authorized distributors, verified agency marks, and the paperwork kept on file	127
13.8	Cheap uncertified wall adapters designed into the product	A certified supply from a real brand with a test report and a certificate number you can look up	128
13.9	Suspiciously cheap "original" ICs from unauthorized sellers	Mouser, Digi-Key, Arrow, Farnell, and other authorized distributors, for every IC on the board	129
13.10	A critical part with one factory on earth, or one already on a countdown	Lifecycle status checked on every part, with a second source or a consciously accepted risk	130

The Six Production Filters

Every entry in this handbook fails at least one of the same six filters, and most fail several. Once you move from a bench prototype to a product you plan to manufacture, run every part, every technology, and every design decision through this list before you commit to it.

1 Reliability

It has to work across the full temperature range, the full supply voltage range, and the full life of the product, not just at room temperature on a fresh battery.

2 Lifecycle

The part must be in active production with a real datasheet, and anything marked not recommended for new designs is living on borrowed time.

3 Compliance

It has to clear RoHS, the certification tests for every market you sell in, and the regional rules that now cover batteries, passwords, and USB connectors.

4 Power

It has to fit inside the power budget, and that matters most in anything running on a battery.

5 Safety

Anything that touches mains voltage, holds a lithium cell, or makes a safety claim needs the isolation, protection, and documentation that hazard calls for.

6 Manufacturability

It has to be sourced from a channel you can trust and built and tested with a repeatable process at the volume you plan, at a cost that still leaves you a margin.

Get a second set of eyes before you spend real money.

Every decision in this handbook is an everyday conversation inside the Hardware Academy, where you can get feedback on your specific schematic, layout, enclosure, and supplier choices from me and a team of engineers who have taken real products into production.

PredictableDesigns.com/academy