SDR Product Handbook: choosing a software-defined receiver or RTSA

Why the landscape looks the way it does, what separates one class of device from the next, and how to choose without being told what to think.

ABSTRACT

There are more software-defined radios on the market than any engineer can evaluate, and from a distance they look interchangeable: similar frequency ranges, similar bandwidths, and prices that cluster by class. They are not interchangeable, and the reason they look alike is more interesting than the differences between them. This handbook explains the landscape from the silicon up. It sets out the three receiver architectures in use, shows why so many unrelated platforms stop at the same frequency and carry the same impairments, works through what an SDR datasheet does and does not tell you, surveys the software world that sits on top, and then draws the line between a device that produces samples and an instrument that produces measurements. It ends with a selection method and a reference section.

WHO THIS PAPER IS FOR

RF and systems engineers choosing hardware for a monitoring, test or research system. **Product architects** deciding whether to buy a module or build a receiver. **Researchers and teaching staff** who need to know what a result from a given class of device is worth. **Buyers and specifiers** who have to write a requirement that several vendors can answer honestly.

CONTENTS

1 How to use this handbook

This is a reference document about a market, not a catalog. It is arranged so that the physics comes before the products and the products come before the recommendations, because the order matters: almost every bad SDR purchase this handbook is meant to prevent came from comparing two datasheets without knowing what the numbers on them were free to mean.

Table 1 is the router. Find the row that matches why you opened this, and start where it says.

Table 1. A router, not a summary. The handbook can be entered at any part; the parts are ordered so that each one earns the vocabulary the next one uses.

If you are	Start at	Why
Trying to understand why every SDR seems to stop at 6 GHz	Part I, section 5	It is one component, and the answer explains most of the landscape
Reading two SDR datasheets and unable to tell them apart	Part II, section 16	The fields that differ are rarely the fields that matter
Choosing a framework or an abstraction layer	Part III	The software choice outlives the hardware choice
Deciding whether you need an instrument or a receiver	Part IV, section 32	This is the only question in the handbook with a wrong answer
Writing a purchase specification	Section 50 and the tables after it	Requirements that several vendors can answer honestly
Looking for a number	The quick reference, section 62	Everything with its condition attached

What this handbook will not do

It will not tell you that one manufacturer's radio is worse than another's. That is partly a matter of policy, since a quantitative comparison between named products is a claim somebody has to stand behind in front of a customer, and partly because it would be bad reference writing. The devices in this market are not ranked along one axis. They occupy different points on several, and a device that is wrong for one job is the obvious answer to another.

What it will do is name the axes, say where the boundaries fall and why, and cite a source for every external fact so that a reader can check it. Where a Berkeley Nucleonics product is the answer, this handbook says so and shows the arithmetic. Where it is not, it says that too, and section 45 exists for exactly that purpose.

> HOW SPECIFICATIONS ARE HANDLED HERE
>
> Every Berkeley Nucleonics figure in this handbook is taken from a published datasheet and carries its measurement condition. Every external figure is attributed to its source in the text and listed in section 66. Anything that is documented rather than published, or that this handbook infers rather than cites, is marked and collected in the verification note at the end. **A specification without its condition is not a specification**, and that rule is applied to other people's numbers as strictly as to ours.

PART I

Why the landscape looks the way it does

Three architectures, one dominant component, and the impairments that follow from them. No product is named in this part.

2 What software-defined actually moved

A radio has always been a chain: capture, select, amplify, convert, decide. What changed over the last three decades is not the chain but the position of one line drawn across it, the line between what is fixed in hardware at manufacture and what can be changed later in software.

In a fully analog receiver that line falls at the far right. The demodulator is a circuit, the filter is a component, and changing the modulation means changing the radio. In a software-defined receiver the line has moved left, past the demodulator, past the filtering, in some designs all the way to the antenna. Everything to the right of the line is code.

Two things pushed it there. Converters became fast enough to digitize a useful span directly, and general-purpose processors became fast enough to do the arithmetic that used to need dedicated silicon. Neither was a radio-engineering advance. The radio world inherited them.

> THE CONSEQUENCE WORTH HOLDING ON TO
>
> **Moving the line left does not improve the radio. It moves the responsibility.** Every function that becomes software becomes yours: yours to write, yours to validate, yours to maintain across a hardware refresh. A device that hands you more of the chain is offering you more freedom and more work in the same gesture, and the whole art of choosing in this market is deciding how much of that trade you actually want.

3 How the market got here, and where it is going

Three shifts made this market, and knowing them explains why its products cluster the way they do and what is likely to change next.

- **Converters crossed a threshold.** Once an analog-to-digital converter could sample fast enough and with enough dynamic range to digitize a useful span, the filtering and demodulation that had been circuits became arithmetic. This is what made a general-purpose receiver possible at all.
- **Integration collapsed the cost.** A complete receive and transmit chain on one part, with a published reference design around it, turned building a radio into an integration exercise. That is the single largest reason there are so many products and so little difference between many of them.
- **The software became the asset.** Open frameworks and abstraction layers accumulated a body of validated processing that no single vendor could match, which moved the durable value from the hardware to the code and made hardware portability a purchasing requirement rather than a nicety.

Two directions follow from those, and both are visible now rather than speculative.

- **Interoperability is becoming a specification.** The DIFI work described in section 23 exists because digitizing an intermediate frequency destroyed the natural interoperability that an analog IF had, and buyers noticed. Expect transport format to appear on requirement lists where it used to appear on nobody's.
- **The interesting frequencies are above the integrated-transceiver ceiling.** Millimeter wave work, satellite ground segments and modern radar all live where a single integrated part cannot reach, which is where the architectural distinction in section 5 becomes a purchasing distinction.

> WHAT HAS NOT CHANGED, AND PROBABLY WILL NOT
>
> **Calibration remains a property of hardware.** Every one of the shifts above moved value into software, and none of them moved the amplitude reference there. A correction can be applied in code, but it has to have been measured against a standard by somebody with the equipment to do it, and that step has resisted every attempt to commoditize it. It is the reason the boundary in Part IV has stayed where it is while everything around it moved.

4 Three architectures, and what each one gives up

Almost every receiver in this market is one of three things. The differences between them are not a matter of quality; each solves the conversion problem differently and pays for it somewhere else, and knowing which one you are holding explains most of its behavior.

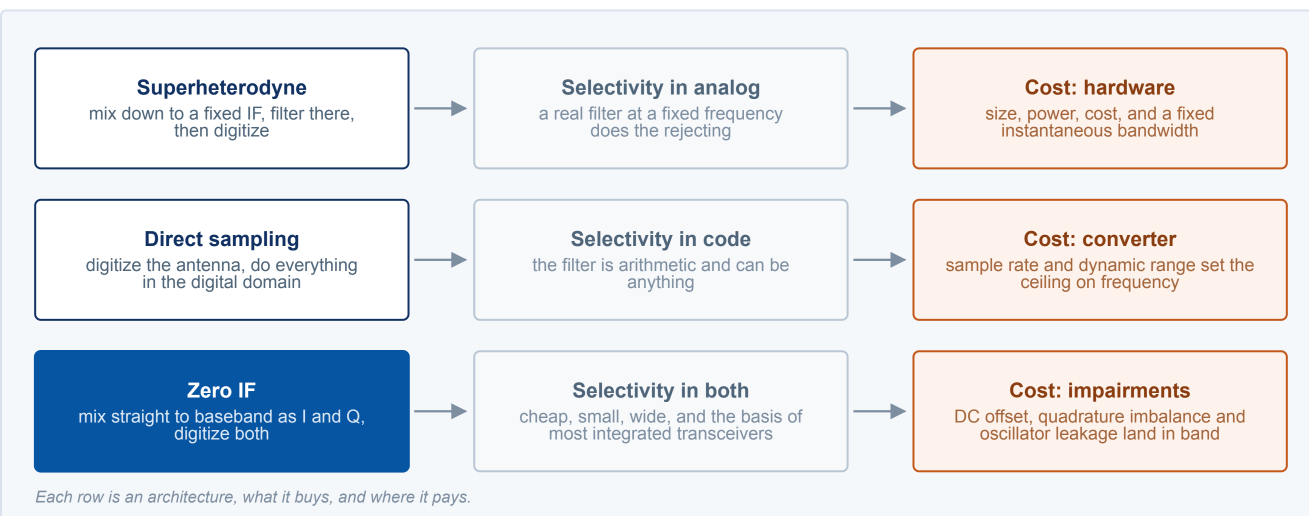

Figure 1. The three receiver architectures in current use, with the price each one pays. Read the third column as the thing you will spend your time on. A superheterodyne receiver costs money and space; a direct-sampling receiver costs converter performance and caps your frequency; a zero-IF receiver costs you three specific impairments that land inside the band you are trying to measure. **The third row is the one to know best**, because it is the architecture behind most of the products in this market. SCHEMATIC

Figure 1 understates one thing. These are not exclusive: a real instrument often uses a superheterodyne front end to reach high frequency and a zero-IF or digital down-conversion stage behind it, so the question is rarely which architecture a device is but which one is doing the work at the frequency you care about.

5 The silicon underneath, and why platforms cluster

Here is the single most useful fact about this market, and the one that explains the most at the least cost. **A large share of the software-defined radios on sale are built around the same family of integrated transceiver.** A second cluster has grown up beside it on a different kind of silicon, and section 6 takes that one; this section takes the cluster that still accounts for most of the units.

Analog Devices publishes the AD9361 as a 2 by 2 agile transceiver whose receive local oscillator covers 70 MHz to 6.0 GHz and whose transmit oscillator covers 47 MHz to 6.0 GHz. The AD9364 is the 1 by 1 variant, and the two share a register map and a package, so a design can migrate between them. Analog Devices also publishes a complete open reference design around the family, with schematics, gateware, firmware, embedded Linux and host drivers.

> **WHY THAT ONE PARAGRAPH EXPLAINS SO MUCH**
>
> If most platforms are built on one transceiver family, then most platforms inherit its frequency limits, its architecture and its impairments. **That is why so many unrelated products stop at 6 GHz**, why their noise and linearity figures cluster inside a narrow band, and why the same three zero-IF artifacts appear in the same places on all of them. The differences between such products are real, but they are differences in what was built around the same core: the clock, the front end, the filtering, the transport, the gateware and the software. Read a datasheet in that light and it becomes much easier to see what a vendor actually contributed.

Two consequences follow, and they point in opposite directions. The good one is that a huge amount of open work, driver support, framework integration and community knowledge accumulates around a common part, which is why an engineer can get a signal on screen in an afternoon. The awkward one is that a specification shared by a dozen products is not a differentiator, and a purchase decision made on it is a decision made on nothing.

> **WHERE THE 6 GHZ CEILING COMES FROM, AND WHAT IT IS NOT**
>
> It is a property of a component, not a law. Above it, a receiver needs a different front end: a superheterodyne conversion chain, more filtering, and a great deal more engineering in the microwave path. That is why the market thins out above 6 GHz and why the products that go there are a different class of thing, at a different price, with a different kind of specification behind them.

6 The second cluster, and where the digitizer moved

The account above is true and it is incomplete, because it describes a part that first shipped in 2013. A second cluster has formed since, and it is built on the opposite idea. Rather than mixing a signal down to baseband and digitizing it there, **it digitizes the radio frequency directly** and does the mixing in arithmetic.

The representative parts are radio-frequency systems-on-chip: converter tiles, programmable logic and processor cores on one die. AMD publishes the Zynq UltraScale+ RFSoC family with sampling rates reaching several gigasamples per second, its digital-front-end variant carrying the direct-sampling input range to 7.125 GHz. **And this cluster is moving fast.** In December 2024 AMD announced the Versal RF series, which it describes as the fifth generation of its direct-RF devices, with direct sampling quoted to 18 GHz at up to 32 gigasamples per second, a pairing worth reading with section 48 in hand, since 32 GSPS puts the first Nyquist zone at 16 GHz, so the top of that input range is deliberate higher-zone operation rather than first-zone sampling. The practical effect is that the boundary drawn in section 7 keeps moving outward: the mixer, the channelizer and the filtering that used to be analog become logic, and the frequency at which that stops being possible rises with each generation.

Table 2. The two clusters. **Neither is the successor of the other**, which is the point most often missed: one moved the digitizer to save cost and power, the other moved it to buy bandwidth, and the impairments swap rather than disappear. A team that has learned to correct quadrature imbalance meets interleaving spurs instead.

	Integrated zero-IF transceiver	Direct RF sampling system-on-chip
How the signal reaches the converter	Mixed down to baseband, then digitized at baseband	Digitized at radio frequency, then mixed down in logic
Typical reach	70 MHz to 6.0 GHz	Published direct-sampling inputs from about 7 GHz on shipping parts to 18 GHz on the newest announced generation, the top of which is higher-Nyquist-zone operation at the quoted converter rates rather than first-zone sampling
Instantaneous bandwidth	Tens of MHz, at most a few hundred	Hundreds of MHz to several GHz
Characteristic impairments	DC offset, quadrature imbalance, LO leakage, all correctable and all recurring	No image and no quadrature error to correct, but converter linearity, clock jitter and interleaving spurs dominate instead
What it costs	Low unit cost, low power, small board	High power, serious thermal design, and gateware skill as a permanent staffing requirement
Who it suits	Volume products, development, anything power-constrained	Wideband instruments, phased arrays, base-station and radar development

WHAT THE SECOND CLUSTER DOES NOT CHANGE

Neither cluster makes a measurement. A direct-sampling front end has the same gap this handbook returns to in section 32: unless a correction is applied against a traceable standard, a wider, faster, more capable digitizer produces wider, faster, more capable numbers with no absolute meaning. Moving the converter closer to the antenna raises the ceiling on bandwidth. It does nothing to the amplitude axis.

For a buyer the second cluster matters in two ways. It is where the wideband end of the market is going, so a requirement written today against a few tens of megahertz may be unwritable in five years. And it is the reason a development board can now quote a bandwidth figure that used to belong to instruments, which makes the question of what a number is guaranteed against more important than it was, not less.

7 Where the programmable boundary falls, drawn

Draw the line between what is fixed in hardware and what is left to software, because every device in this handbook can be placed on it and the placement predicts most of its behavior.

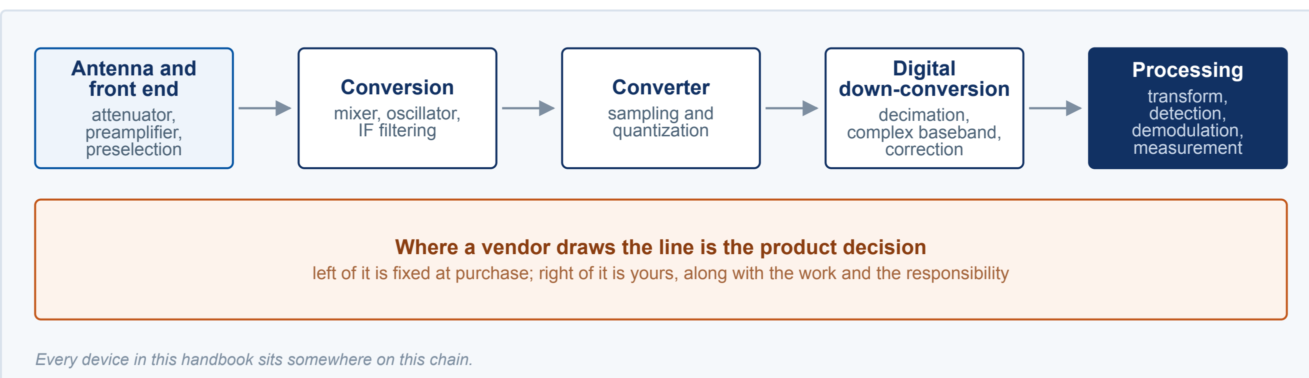

Figure 2. The receiver chain, with the boundary that actually distinguishes products. A consumer dongle draws the line immediately after the converter and hands you everything, uncorrected. A closed benchtop analyzer draws it at the far right and hands you a trace. **The interesting position is between the correction and the processing**, because that is the only place a device can hand you samples that already mean something in dBm. Where a vendor draws this line is the product decision, and reading a datasheet is largely the work of finding out where they drew it. SCHEMATIC

Figure 2 also explains why the phrase software-defined radio has stopped discriminating between products. Nearly everything on sale is software-defined by some reading of the term. The useful question is not whether a device is software-defined but **how far left the line is drawn and what crosses it**, which is a question with a checkable answer.

8 Sampling, and the three numbers people confuse

Three quantities get used interchangeably in this market and they are not the same: sample rate, instantaneous bandwidth and usable bandwidth. Getting them straight removes a large class of purchasing mistakes.

A real band-limited signal needs a sample rate of twice its bandwidth, because each real sample carries one degree of freedom and the spectrum is conjugate symmetric. A complex sample carries two, so a complex sequence at rate f_s represents a full span of width f_s about the tuned center:

$$B_{\text{complex}} = f_s \quad \text{versus} \quad B_{\text{real}} = f_s / 2 \tag{1}$$

where

B is the representable bandwidth, in hertz

f_s is the sample rate, in samples per second

No information is gained by equation (1). B complex samples per second is 2B real numbers per second either way. What changes is the bookkeeping, and almost every SDR datasheet quotes the complex form without saying so.

USABLE BANDWIDTH IS ALWAYS NARROWER THAN THE RATE

Anti-alias and decimation filters need a transition band, so the span you can actually trust is narrower than the span the sample rate represents. **A ratio around 0.8 is typical.** This is why a device advertising 61.44 MSPS is not a 61.44 MHz analyzer, and why quoting a file's sample rate as its bandwidth is the standard avoidable error in this field. Ask for both numbers, and treat a vendor who publishes only one as having answered the easier question.

WHAT AN IQ SAMPLE ACTUALLY IS, SINCE EVERYTHING DOWNSTREAM ASSUMES IT

A real receiver multiplies the incoming signal by two copies of the local oscillator a quarter cycle apart and low-pass filters each result. The two outputs are the in-phase and quadrature components, and **a sample is that pair, not two samples**. The pair carries amplitude and phase together, which is why a complex stream at rate f_s represents a span f_s wide rather than $f_s/2$, and why the spectrum it produces is not symmetric about the center and can distinguish a signal above the tuned frequency from one below it. A real-only stream cannot: it folds the two together, which is the same fact as the complex-Nyquist relation above, seen from the other side.

DIGITAL DOWN-CONVERSION, WHICH IS WHAT THE DECIMATION IS ATTACHED TO

Between the converter and everything downstream sits a stage that does in arithmetic what a mixer does in hardware: it multiplies the samples by a numerically generated oscillator to shift the band of interest to zero frequency, filters away everything else, and only then decimates. **The order matters and is the whole point:** tune, filter, then throw samples away, because discarding first folds the rejected energy back in where it cannot be separated. This is the stage the second silicon cluster uses instead of an analog mixer, it is what a published decimation range is a specification of, and it is why a receiver can present a narrow, low-rate, precisely tuned stream from a converter running far faster and far wider.

DECIMATION, SINCE EVERY LATER RELATION USES IT

A receiver rarely delivers samples at the converter's rate. **Decimation by *D* keeps one sample in *D* and discards the rest**, which divides both the delivered rate and the representable bandwidth by *D*. It is never just discarding: a filter has to remove everything outside the surviving band first, or that energy folds in and cannot be separated afterwards. Done properly it is free resolution, since a narrower band carried at a lower rate is less data for the same information. Every *D* in this handbook is that factor, and the reason it appears in the intercept relation is that a longer transform over a lower rate takes proportionally longer in real time.

9 Dynamic range, and why the bit count misleads

The second field people over-read is converter resolution. An ideal converter quantizing a full-scale sine into *b* bits achieves a signal-to-noise ratio of:

$$\mathrm{SNR} = 6.02b + 1.76\ \mathrm{dB} \qquad (2)$$

where

b is the converter resolution, in bits

Equation (2) is an upper bound nobody reaches. Real converters fall short, and the honest figure is the effective number of bits derived from measured signal-to-noise and distortion:

$$\mathrm{ENOB} = (\mathrm{SINAD} - 1.76)\ /\ 6.02 \qquad (3)$$

where

$SINAD$ is measured signal to noise and distortion ratio, in dB

$ENOB$ is effective number of bits

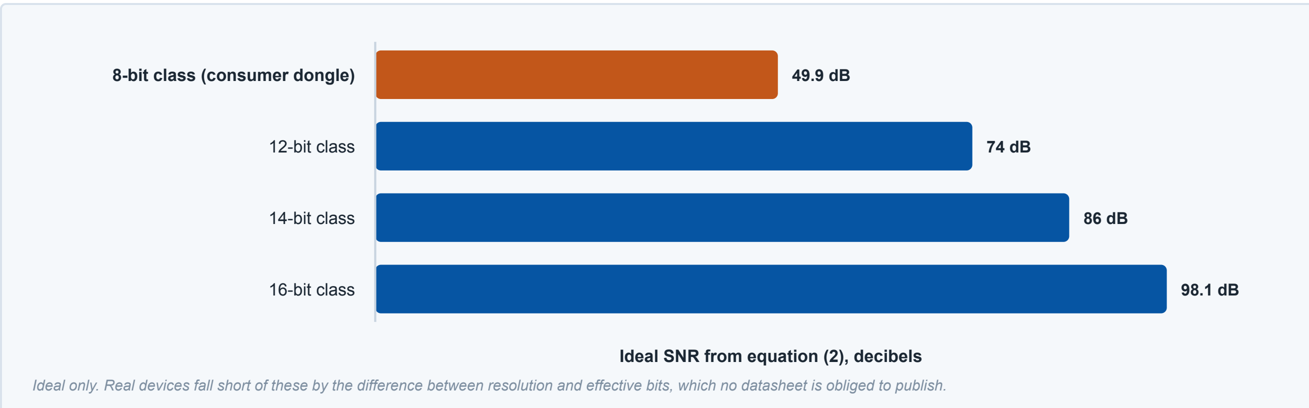

Figure 3. Equation (2) evaluated across the resolutions this market actually ships. The spread from a consumer dongle to an instrument converter is about 48 dB of ideal signal-to-noise, which is the single largest technical gap in the landscape and the one least visible on a specification sheet. **Read these as ceilings, not as performance.** A device is entitled to publish its bit count and say nothing about SINAD, and most do, which means the number on the page is the best case for a part that was never measured in the product. DERIVED

Figure 3 explains a pattern buyers notice and rarely account for. Two devices quoting the same frequency range and the same bandwidth can differ by tens of decibels in what they can see next to a strong signal, and nothing in the headline specification says so. Processing gain recovers some of it, because narrowing the resolution bandwidth by a factor of *k* lowers the noise floor by ten times the log of *k*, but processing gain cannot recover a signal the converter clipped.

10 The two ends of the range, and the span between them

The bit count describes the converter. What a buyer needs is the range of the whole receiver, and that is set at the bottom by noise and at the top by distortion. Four quantities fix it, and a datasheet that omits them has not described the device.

$$\mathrm{MDS} = -174\ \mathrm{dBm/Hz} + \mathrm{NF} + 10\log_{10}(B) \qquad (4)$$

where

MDS is the noise floor referred to the input, in dBm

NF is the receiver noise figure, in decibels

B is the measurement bandwidth, in hertz

The constant in equation (4) is thermal noise power density at the 290 K reference temperature, which is a convention rather than a law: a cryogenically cooled front end sits below it, and a warm antenna looking at a warm scene sits above it. For any receiver operating at room temperature it is the practical floor, so **a sensitivity claim can be checked against it in one line.** A receiver with a 10 dB noise figure measured in a 1 kHz bandwidth cannot hear below −174 + 10 + 30, which is −134 dBm, and a claim of −150 dBm in that bandwidth is not a good product but an arithmetic error or an unstated bandwidth.

Table 3 is the set of four.

Table 3. The four numbers that bound a receiver. **They are not independent.** The gain that improves noise figure worsens intercept, which is the trade of section 14 arriving from the front end, and it is why a preamplifier is a decision rather than an upgrade.

Quantity	What it means	What to ask for

Quantity	What it means	What to ask for
Noise figure	How much the receiver's own noise adds to the thermal floor	The figure with its frequency and its gain setting. It changes with both
1 dB compression point (P1dB)	The input level at which gain has fallen by 1 dB, so the receiver has begun to flatten the signal	Input-referred, at a stated gain. This is the top of the honest range, not the damage level
Third-order intercept (IIP3)	An extrapolated level describing how fast intermodulation products grow. Two strong signals produce products that rise **three decibels for every decibel** of input	Input-referred, with the tone spacing used. A number without its spacing is not checkable
Spurious-free dynamic range (SFDR)	The span between the noise floor and the level at which the first false product appears	The bandwidth it was measured in. SFDR is not a single property of a receiver; it moves with resolution bandwidth

WHY A SINGLE DYNAMIC RANGE FIGURE SHOULD BE REFUSED

The three ranges a receiver has are different numbers: the **display range** between the reference level and the floor, the **instantaneous range** the converter spans in one acquisition, and the **spurious-free range** available before the receiver invents a signal. A datasheet quoting one figure called dynamic range has chosen the most flattering of the three and has not said which. Ask which, and ask in what bandwidth. Two thirds of the difference between IIP3 and the noise floor is the usual construction of the third one, so it can be reconstructed if the parts are published.

11 Phase noise, and the signal you cannot see next to a strong one

Noise and distortion set the two ends of the range. A third mechanism decides what you can see **next to** something strong, and it is the one most often left out of a comparison because it does not appear as a single number.

No oscillator produces one frequency. A real local oscillator produces a carrier with skirts of noise either side, and **phase noise** is the power in those skirts, quoted as decibels relative to the carrier in a one hertz bandwidth at a stated offset. So a figure like −110 dBc/Hz means nothing until you are told the offset it applies at, and a vendor quoting one offset has chosen the flattering one.

RECIPROCAL MIXING, WHICH IS WHY THIS MATTERS TO A RECEIVER RATHER THAN TO A TRANSMITTER

When a strong signal enters a receiver, it mixes with the local oscillator's noise skirts as well as with its carrier. The result is that **the strong signal's energy is smeared across the offsets around it by the receiver's own oscillator**, raising the apparent noise floor near it. The displayed floor beside a strong carrier is roughly the carrier level, plus the phase noise at that offset, plus ten times the log of the resolution bandwidth. A −10 dBm carrier, an oscillator at −110 dBc/Hz at the offset you care about, and a 1 kHz resolution bandwidth put the floor beside it near −90 dBm, **whatever the receiver's own noise figure is.** Processing gain lowers this floor exactly as it lowers the thermal one, 10 dB per decade of resolution bandwidth, so a weak neighbor does become visible in a narrower resolution bandwidth. What narrowing cannot do is change which floor dominates, because both fall at the same rate: a receiver that is phase-noise-limited beside a strong carrier is still phase-noise-limited at every resolution bandwidth, and the only fixes are a better oscillator or attenuating the carrier ahead of the mixer.

Two consequences follow that a datasheet comparison will not show you. A receiver with an excellent noise figure and a mediocre oscillator finds weak signals in a quiet band and loses them next to a strong neighbor, which is the situation every real deployment is in. And **an adjacent-channel measurement is a measurement of your instrument's oscillator until you prove otherwise**, which is why the honest way to report one is with the instrument's phase noise at the offset in question stated alongside it.

12 What zero IF costs, and where it shows

The zero-IF architecture mixes the wanted band straight down to baseband as an in-phase and a quadrature stream. It is elegant, it is cheap, it integrates well, and it has three characteristic impairments that every user of such a device eventually meets. They are worth naming because they are visible on a screen, and once you can name them you stop misdiagnosing them as signals.

- **A DC offset at the center of the span.** Any offset in the baseband path appears at zero frequency, which after conversion sits exactly at the tuned center. It looks like a carrier and it is not one. Tuning deliberately off the signal of interest, which is what a local oscillator offset does, moves the wanted signal away from it.
- **Quadrature imbalance, seen as an image.** If the I and Q paths differ in gain or are not exactly ninety degrees apart, a signal at one offset produces a mirror at the negative of that offset. The mirror is an artifact of the receiver and moves when you retune, which is the test that distinguishes it from a real emitter.
- **Local oscillator leakage.** The oscillator is at the tuned frequency by construction, so any leakage back to the input reappears at the center of the span, reinforcing the same spot the DC offset occupies.

THE DIAGNOSTIC THAT SETTLES ALL THREE

Retune by a known amount and see what moves. A real signal keeps its absolute frequency and therefore shifts in the display. An artifact of the receiver stays where it was relative to the center, because it is generated relative to the center. That one habit resolves more spurious-signal reports than any amount of specification reading, and it costs ten seconds.

13 The five settings that decide what a screen shows

Every instrument capture in this handbook was taken with five controls set, and every one of them changes the answer. A reader who has only ever used a receiver through a flowgraph will not have met them, and a reader who has used an analyzer for years may never have been told what they do. They are the difference between two people pointing the same instrument at the same signal and reporting different numbers.

Table 4 is the set.

Table 4. The five display controls, plus a sixth that sits outside the list and matters as much: **the reference level and the input attenuation it sets**, which decide where the instrument's own range is placed and therefore whether a signal is compressed or buried. **Any amplitude number you report is a statement about all six**, which is why the reproducibility checklist in section 65 asks for every one of them and why a screenshot without them is a picture rather than a record.

Control	What it does	What goes wrong when it is wrong
Resolution bandwidth (RBW)	The width of the filter, or equivalently of the transform bin, that each displayed point represents	It sets the noise floor. Ten times narrower is 10 dB lower, by the relation in section 14. Two traces at different RBW are not comparable, and a noise floor quoted without its RBW means nothing
Video bandwidth (VBW)	Smoothing applied after detection, on the envelope rather than on the signal	Narrow VBW makes a trace look clean and pulls down the peaks of anything short. Set VBW at least three times RBW when measuring pulses, and well below RBW only when you are deliberately averaging noise
Detector	How many samples inside one displayed point are reduced to the one value drawn: peak, sample, average or RMS	**Peak overstates noise; sample understates bursts; only an RMS detector, or an average detector that averages in power rather than in log units, gives a correct power for noise-like signals, and a log-scale average under-reads Gaussian noise by 2.51 dB.** A channel power measured on a peak detector is wrong and looks fine
Window function	The taper applied before the transform, which trades frequency resolution against leakage into neighboring bins	It changes the effective noise bandwidth, so the same transform size gives a different RBW under a different window. A flat-top window measures amplitude well and resolves poorly; the reverse holds for narrower windows
Trace mode	What happens to a point across successive sweeps: clear write, max hold, min hold or average	Clear write shows one instant, max hold shows the union of everything seen. Reading a max-hold trace as though it were an instant, or the reverse, is the single most common misreading of a spectrum display

> **THE RELATION BETWEEN A TRANSFORM AND A RESOLUTION BANDWIDTH**
>
> A transform of N points across a sample rate f_s produces bins spaced f_s/N apart, but the resolution bandwidth is wider than the spacing, by a factor that belongs to the window. That factor is the window's equivalent noise bandwidth. For the Blackman-Nuttall window it is 1.9761, so 30.518 kHz bins are a 60.31 kHz resolution bandwidth, which is what the display reports. **Bin spacing is geometry; resolution bandwidth is the number you quote.**

14 Processing gain, and what it can and cannot recover

One more relation belongs in Part I because it is used to justify purchases it does not support. Narrowing the resolution bandwidth lowers the displayed noise floor:

$$\Delta N = 10 \log_{10}(k) \quad (5)$$

where

k is the factor by which resolution bandwidth is reduced

ΔN is the reduction in displayed noise floor, in decibels

Equation (5) is why a signal invisible in a 1 MHz resolution bandwidth appears cleanly in 1 kHz: a factor of a thousand is 30 dB of noise floor. It is real, it is free, and it is the reason a narrow-resolution sweep finds things a wide one does not.

> **THE THREE THINGS PROCESSING GAIN DOES NOT DO**
>
> **It does not recover a clipped signal.** Once a strong input has driven the converter into its limit, the information is gone and no amount of averaging brings it back. **It does not shorten the guaranteed intercept.** A narrower resolution bandwidth means a longer transform, and therefore a longer intercept guarantee, so the very setting that finds a weak continuous signal is the setting that misses a short burst. **And it does not create dynamic range against a strong neighbor**, because the spurious products a strong signal generates inside the receiver fall inside the narrow bandwidth too.

That second point is the trade of section 33 arriving from the other direction, and it is worth stating as a rule. **Sensitivity and intercept pull against each other through the same product**, so a specification that demands both an extremely low noise floor and a microsecond intercept guarantee is asking for two settings at once. Decide which the application actually needs, or specify two configurations and accept switching between them.

15 Noise, gain and the front end you attach

A receiver's published noise figure describes the receiver. What decides whether you hear anything is the whole chain from antenna to converter, and the first stage dominates it. The cascade relation settles most front-end arguments on its own:

$$F_{\text{total}} = F_1 + (F_2 - 1)/G_1 + (F_3 - 1)/(G_1 G_2) + \ldots \qquad (6)$$

where

F is the noise factor of each stage, linear rather than in decibels

G is available gain of each stage, linear

> **WHAT EQUATION (6) ACTUALLY TELLS A BUYER**
>
> **The first stage sets the noise figure, and everything after it matters less in proportion to the gain in front of it.** A low-noise preamplifier at the antenna improves a mediocre receiver far more than a better receiver improves a bad front end. It also costs you linearity, because every strong signal in the band is now amplified too, which is why the preamplifier that helps a weak-signal survey ruins a measurement next to a transmitter.

Worked once, the relation stops being an argument. Take a receiver with a 15 dB noise figure and three decibels of cable ahead of it. A passive loss contributes its own loss as noise figure, so the chain comes to 18.0 dB and the cable cost exactly what it lost. Now put a 1.0 dB, 20 dB-gain preamplifier at the antenna, ahead of the same cable and the same receiver. The relation gives 1.259 plus 0.010 plus 0.611, which is 1.880, and ten times its log is **2.74 dB**. The receiver did not change and the chain improved by 15.3 dB, all of it bought by twenty decibels of gain standing in front of everything else. **The bill arrives on the other axis:** every signal reaching the receiver is now 20 dB stronger, so the compression and intercept points of the whole chain drop by about the same 20 dB, and a preamplifier that rescues a quiet rural site will drive the same receiver into distortion within sight of a broadcast mast.

Three practical consequences follow, and they apply whatever class of device you buy.

- **Cable loss before the first amplifier is noise figure**, decibel for decibel. Several decibels per meter is normal in flexible assemblies at millimeter-wave frequencies, so a long run to a mast-top antenna can cost more sensitivity than any receiver choice will recover. Amplify at the antenna or accept the loss knowingly.
- **Choose the antenna before the receiver.** A broadband discone gives up roughly 6 dB against a log-periodic and considerably more against a horn on the same band, and the antenna factor has to be entered before any field-strength number means anything at all.
- **Preselection is what lets a wideband receiver work in a real environment.** A device that digitizes a wide span sees every emitter in that span at once, so the strongest one sets the gain and everything weak lives in whatever dynamic range is left.

PART II

The product landscape

What the classes of device actually are, what a datasheet does and does not tell you, and the four fields that decide most purchases.

16 Reading an SDR datasheet

An SDR datasheet is a shorter document than an instrument datasheet, and the difference is not laziness. Most of the fields an instrument publishes describe accuracy, and a device that makes no accuracy claim has nothing to put in them. Knowing which fields are missing is more useful than reading the ones that are there.

Table 5 is the field list, with what each one is worth.

Table 5. The fields, and what each is silent about. **Read the third column first.** A datasheet is a description of what a vendor is prepared to be held to, and the shape of what is missing tells you what class of device you are looking at more reliably than any number that is present.

Field	What it tells you	What it does not
Frequency range	Where the tuner can be pointed	Whether performance is flat across it. It rarely is
Instantaneous bandwidth	How much spectrum is digitized at once	Whether the whole of it is usable. Filters need a transition band, so the usable span is always narrower than the number quoted
Sample rate	How fast samples leave the converter	**Not the same as bandwidth.** Quoting a sample rate as a span is the commonest error in this field, and many datasheets invite it
ADC bits	The converter's nominal resolution	The dynamic range you get. Effective bits after noise and distortion are always fewer, and are what SINAD would tell you if it were published
Sustained versus burst rate	Whether the number is a stream or a buffer fill	Anything, if the datasheet quotes only one figure. A device that publishes both is telling you something about its honesty
Transmit power	The output at some frequency and some setting	Linearity, spectral purity or whether the figure holds across the range
Reference oscillator	Frequency stability against a disciplined source	Holdover, aging or temperature coefficient, which is what actually matters once the discipline is removed
Absolute amplitude accuracy	**Usually nothing, because it is usually absent**	This is the field whose absence defines the category. See section 32

17 The classes of device

Products in this market fall into four groups. The grouping is by where the boundary between fixed hardware and user responsibility falls, which is the distinction from Part I applied commercially.

Named examples make the classes concrete, so here are the ones a newcomer actually meets. **The first row** is where most people start: receivers built on the RTL2832U demodulator, sold as RTL-SDR sticks, and the Airspy and SDRplay RSP families a step above them. **The second row** is the research bench: the ADALM-PLUTO, the USRP B-series, LimeSDR and bladeRF, most of them built on the transceiver family of section 5, alongside the HackRF as a wideband half-duplex outlier. Evaluation boards for the direct-sampling parts of section 6 sit at the top of this row at a very different price. **The third row** is rarely famous, because a module that ends up inside somebody else's product is sold on supply life and interface stability rather than on brand. **The fourth row** is where the ICX line sits, and it is a different category rather than a better second row.

These are named as landmarks, not as recommendations, and nothing here compares them against each other or against ours. Each publishes its own specifications; read them against the field list in section 16 and the class becomes obvious from what is missing.

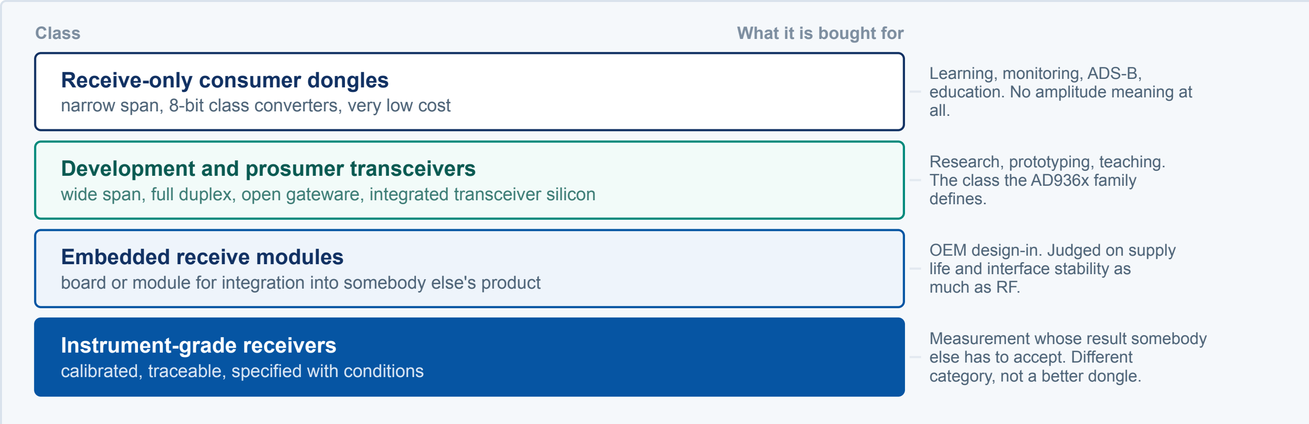

Figure 4. The four classes, and the job each is bought to do. The boundary that matters is between the third row and the fourth, and it is not a matter of degree: everything in the first three rows hands you a number that means something relative to full scale, and the fourth hands you a number that means something in dBm with a stated bound. A device from any row can be the right answer; a device from the wrong row is expensive whichever direction the mistake runs. SCHEMATIC

Read Figure 4 in both directions. Buying up the stack when the job did not need it wastes money on accuracy nobody will use. Buying down it, which is far more common, produces a result that cannot be defended when somebody asks how it was calibrated, and the cost lands months later when the campaign has to be repeated.

One more thing follows from Figure 4 that is easy to miss. The classes are not a price ladder with the same product at four qualities. They differ in what the vendor is prepared to be held to, which is why a device three rows down can outperform one two rows up on a headline number and still be the wrong purchase. **Buy the row, then the model.**

18 Transmit, duplex and multiple channels

Three capability questions separate otherwise similar devices, and each has a trap in it.

- **Transmit is not the mirror of receive.** A device that receives well may transmit poorly, and transmit specifications are usually thinner: an output power figure at one frequency, and often nothing about linearity or spurious content. If the transmitted signal has to be clean rather than merely present, that is a question to ask before purchase.
- **Half duplex and full duplex are different products.** A half-duplex device transmits or receives, not both, which rules out any application that has to hear itself or hear a reply during transmission. The datasheet word to look for is the one that is often missing.
- **Two receivers are not a coherent pair.** Two channels that happen to be on one board are not phase coherent unless they share a reference and their relative phase is specified. Direction finding, beamforming and time-difference work all need the second property, and it is the one least often published.

> **THE QUESTION THAT SEPARATES A CHANNEL COUNT FROM A CAPABILITY**
>
> Ask what is specified about the **relationship** between the channels, not about the channels. Common reference, common clock, hardware trigger, and a stated bound on relative phase or delay. A device with four channels and nothing published about their relationship is a device with four independent receivers, which is a different and much less useful thing than a coherent array.

19 Antennas, preselection and the part nobody budgets for

A receiver is half a system. The other half is what is in front of it, and it is routinely specified last, bought cheapest and blamed least when the numbers disappoint.

Table 6 is the antenna decision in the order it actually matters.

Table 6. The antenna decision. Answer the right-hand column and the other three follow. Note the last row's cost: an active antenna improves sensitivity and reduces the strongest signal a system tolerates, which is exactly the wrong trade in a dense urban environment and exactly the right one on a quiet rural site.

Antenna type	What it buys	What it costs	Right when

Antenna type	What it buys	What it costs	Right when
Broadband omnidirectional (discone, biconical)	Coverage of a wide band from one element, no pointing	Several decibels of gain against a directional element, and no spatial rejection	Survey and monitoring where you do not yet know what is there
Log-periodic	Directivity across a wide band, moderate gain	Pointing, and a physically large element at low frequencies	Direction finding and EMC work across decades of frequency
Horn	High gain and clean pattern over a narrower band	Bandwidth, size and cost	Millimeter-wave work and any measurement wanting the best signal-to-noise on one band
Active antenna with integrated amplifier	Noise figure set at the antenna rather than after a cable run	Dynamic range, since strong signals are amplified too	Long feeder runs and weak-signal work in a quiet environment

ANTENNA FACTOR IS NOT OPTIONAL ARITHMETIC

A receiver measures power at a connector. A regulation, a safety limit or an interference complaint is written in field strength. Converting between them needs the antenna factor in dB per meter, and the conversion is **field strength in dBμV/m = receiver reading in dBm + 107 + antenna factor**, where the 107 assumes a 50 ohm system. A field-strength number produced without that step is wrong by whatever the antenna factor happened to be, which at these frequencies is routinely twenty to forty decibels.

Preselection is the other half of the same subject. A wideband receiver sees every emitter in its span at once, so the strongest one sets the gain and everything weak lives in what is left. In a broadcast-heavy environment a filter ahead of the receiver can buy more usable dynamic range than any change of receiver, and it is cheaper. The trade is that a filter you install is a band you can no longer see, which is why the question belongs in the requirement rather than in the field.

20 What each class actually costs you to own

Table 7 is the part of a purchase that no datasheet covers: what the device asks of you after it arrives. The rows are the classes from Figure 4, and the columns are the costs that show up in a project plan rather than on an invoice.

Table 7. The ownership cost by class. **The third column is where projects actually differ.** Two devices can reach a spectrum display in the same afternoon and be months apart on a number somebody outside the team will accept, and that gap is invisible on both datasheets. Note the fourth row's fourth column: a calibration interval is a cost, but it is a known and scheduled one, which is a different category of problem from a correction table that drifts on nobody's calendar.

Class	You supply	Time to first useful result	What ages badly
Consumer dongle	Every correction, and a span that limits what can be seen at once	An afternoon	Nothing much. It is cheap enough to replace
Development transceiver	The amplitude reference and its maintenance across gain states	A week to a signal, months to a defensible number	The correction table, which drifts and has to be rebuilt
Embedded module	Integration, thermal design, and conformity of the finished product	A quarter, dominated by mechanical and regulatory work	Supply life and interface stability, which outrank RF performance
Instrument-grade receiver	The processing you wanted to write anyway	A day to a defensible number	The calibration interval, which is a scheduled cost rather than a surprise

THE MISTAKE THIS TABLE EXISTS TO PREVENT

Buying down the stack to save money on a project whose output has to be defended. The saving is real and immediate; the cost arrives when a reviewer, an auditor or a customer asks how the amplitude was established, and the honest answer is that it was not.

Nothing recovers that except repeating the campaign with a device that could answer the question.

21 Timing, reference and synchronization

Timing is where distributed systems succeed or fail, and it is consistently the weakest section of an SDR datasheet. Three quantities matter and they are frequently confused.

AND ONE ASSUMPTION UNDERNEATH ALL THREE THAT HAS STOPPED BEING SAFE

Almost every disciplined reference in this market is disciplined by satellite navigation, and **satellite navigation is now a routinely attacked input.** Jamming and spoofing around conflict zones and busy airspace have grown from a curiosity into a daily operational fact, reported in volume by civil aviation. A spoofed receiver does not report an error; it reports a confident and wrong time, which is the failure mode this handbook keeps returning to in a different costume. Two consequences follow for a distributed system: **holdover stops being a specification detail and becomes the survival requirement**, because it is what you run on when discipline is lost or should be rejected, and a monitoring network is one of the few things well placed to detect the attack, since interference in the navigation bands is exactly what a wideband receiver sees. Ask what a node does when its time source disagrees with its own oscillator, and treat a supplier without an answer as having assumed the problem away.

Table 8 separates them, because a datasheet that answers one is often read as having answered all three.

Table 8. The three timing quantities, which are not interchangeable. A parts-per-million figure describes accuracy against a reference that is present; it says nothing about what happens when the reference goes away, and holdover is what a distributed system actually depends on. **Ask for holdover separately or measure it**, because a datasheet that quotes only the disciplined figure has answered an easier question than the one you asked.

Quantity	What it governs	Typical grades
Reference accuracy	How far the tuned frequency is from where you asked	Parts per million against a disciplined source
Holdover	How far it drifts once discipline is removed	Governed by aging and temperature coefficient, and often unpublished
Timestamp accuracy	How well a sample can be placed on an absolute time base	Nanoseconds against a pulse-per-second input

Memorize one conversion, because it turns a timing specification into a distance and makes the consequences obvious. Light travels 0.2998 m in a nanosecond, so an arrival-time uncertainty of a hundred nanoseconds is thirty meters of range uncertainty before any geometry is applied. BNC-AN-110 works that arithmetic through for a distributed radio-location system and finds that the receiver's own contribution is not the term that decides the answer.

22 Form factor, environment and the things that break

A receiver spends its life somewhere, and where decides more about the purchase than most specification sheets acknowledge. Three questions separate a device that works in a laboratory from one that works where the signal is.

- **Power and thermal.** A module drawing ten to twenty watts inside a sealed enclosure is a thermal problem before it is an electrical one, and enclosure temperature moves amplitude through the receiver's temperature coefficient. Published operating classes matter: a base class covering 0 to 50 °C is a laboratory instrument, and a class reaching −40 to +65 °C is a different product even if the RF specification is identical.

- **Ingress and mechanical protection.** A sealed chassis is the difference between a device that survives a wet mast and one that survives a demonstration. Read the specification row rather than the prose: ingress ratings are frequently described in marketing copy and absent from the table.
- **Damage limits.** Every receiver has a maximum input beyond which something stops being reversible. A device with a preamplifier switched in usually has a much lower limit than the same device without, and a field operator who forgets that is the commonest cause of a receiver going back to the factory.

THE DAMAGE LIMIT IS A SPECIFICATION, AND IT IS OFTEN THE ONE NOBODY READS

The ICX-FieldHawk publishes maximum continuous-wave input as +23 dBm at 50 MHz and above with the preamplifier off, falling to +10 dBm below 50 MHz or with the preamplifier on. **That is a 13 dB difference created by a front-panel setting**, and it is the kind of detail worth putting on a laminated card next to any receiver that goes into the field. Ask every vendor for both numbers.

23 Interfaces, transport and the standards that matter

How samples leave the device is a system decision disguised as a connector choice.

- **USB** is convenient, ubiquitous and shares its bandwidth with everything else on the controller. It is the reason a great many overflow problems are diagnosed as software faults when they are transport faults.
- **Ethernet** decouples the receiver from the host and lets it sit where the antenna is, at the cost of needing the network to behave. For distributed sensing this is usually the right answer.
- **PCIe** gives the most throughput and the least flexibility about where the receiver is.

Above the connector there is a question of format, and this is where the market has done something genuinely useful. **VITA 49** is an ANSI standard for carrying digitized RF with its metadata, widely used in satellite and software-defined radio work. Its flexibility is also its problem: two compliant implementations can still fail to interoperate because the standard defines what is possible rather than what is required.

The **DIFI Consortium** was formed under the IEEE Industry Standards and Technology Organization to fix exactly that. Its Digital IF Interoperability Standard, published as IEEE-ISTO Std 4900-2021, is a constrained profile of VITA 49.2: a specific schema that removes the ambiguity, so that where VITA 49 defines what is possible, DIFI defines what is required. The consortium's stated purpose is to replace the natural interoperability that analog intermediate-frequency signals used to have, and to prevent vendor lock-in.

WHY A BUYER SHOULD CARE ABOUT A TRANSPORT STANDARD

Because it is the difference between a receiver you can replace and a receiver you cannot. An analog intermediate frequency was interoperable by physics: any two boxes agreed about a voltage on a cable. Digitizing that link removed the agreement and put a proprietary format in its place, which is how a ground station ends up unable to change one vendor without changing all of them. **If your architecture is distributed and expected to outlive its current hardware, ask about DIFI by name.**

One practical note on asking. The standard has been revised since its 2021 designation, so a supplier who answers "we support DIFI" has not answered the question. Ask which version, and ask whether they have run the consortium's certification material against it, because a packet format is only interoperable to the extent that two implementations have actually been made to talk to each other.

WHERE THE ICX-FIELDHAWK STANDS ON THIS, STATED PLAINLY

The ICX-FieldHawk does not publish VITA 49 or DIFI support. The claim appears on a preliminary portfolio page and on no specification table, so this handbook does not make it. What is published is a documented interface set, SCPI as standard, and language bindings across C, C++, C#, Python, MATLAB, Qt and LabVIEW. A SoapySDR driver module exists and is documented with worked examples, but it is documented rather than specified, so this handbook treats it as demonstrated and not as a published figure. For a distributed ground-segment architecture built on DIFI specifically, the absence of DIFI is a gap worth raising with us before it becomes a surprise.

PART III

The software world

The frameworks, the abstraction layers and the applications. The software choice usually outlives the hardware choice, and it is made with less care.

24 Frameworks, abstractions and applications

Software for this market divides into three layers you choose between, resting on a fourth you do not, and confusing them is the source of most bad architecture decisions. A framework is where you build signal processing. An abstraction layer is how any framework reaches any radio. An application is a finished tool somebody else built.

The names are better stated than talked around, because a reader cannot evaluate an architecture described only in categories. **GNU Radio** is the dominant open framework: a scheduler, a block library and a graphical flowgraph editor, with a long history and an annual conference. **SoapySDR** is the dominant open abstraction layer: a small C and C++ API with language bindings, against which a driver module is written per device family, so a flowgraph written for one radio runs on another. Around them sit finished applications: general-coverage receivers such as GQRX, SDR# and SDRangel, and single-purpose decoders such as dump1090 for aircraft transponders, which are what most people meet first. Vendor software development kits occupy the same layer as the abstraction layer, with the opposite trade: everything the hardware can do, for exactly one vendor.

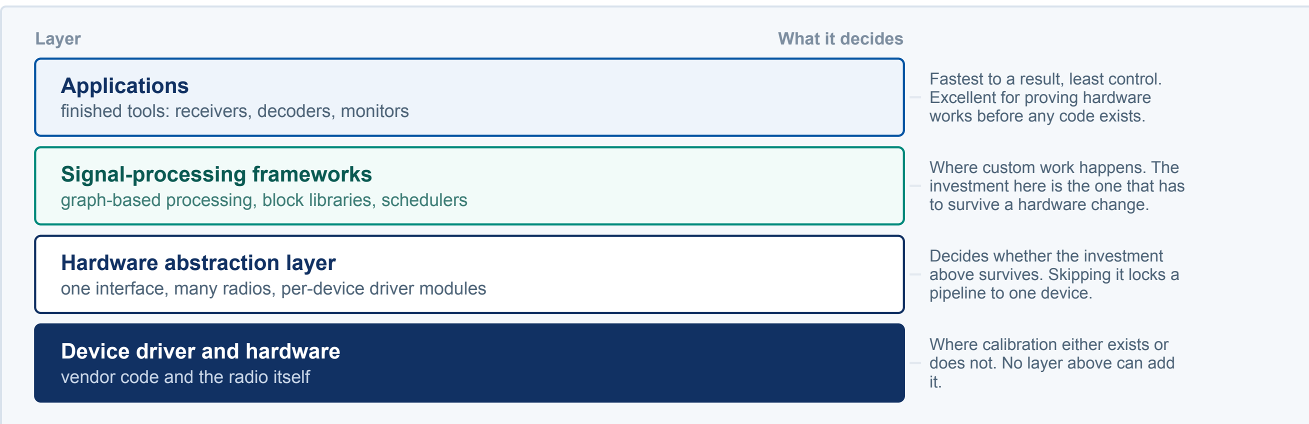

Figure 5. The software stack, and what each layer decides. **Read the second and third rows together.** The expensive, slow, carefully validated work lives in the framework layer, and whether that work survives a change of radio is decided entirely by whether an abstraction layer sits underneath it. A pipeline written straight against a vendor driver is a pipeline with one supplier for its life, and that decision is usually made in an afternoon by somebody who did not know they were making it. SCHEMATIC

Figure 5 carries a warning worth stating directly. The bottom row is the only place calibration can come from. An application cannot add it, a framework cannot add it, and an abstraction layer explicitly does not carry it. If the hardware does not produce corrected samples, no amount of software above it will produce a calibrated measurement.

25 Abstraction layers, and what they deliberately omit

The third layer of the software stack is the one that decides whether the work in the second survives a change of hardware, and it is chosen with the least deliberation of the three. It deserves more, because the decision is close to irreversible once a body of processing has been written against it.

The problem it solves is narrow and expensive. Every vendor ships a programming interface, and each answers the same first questions differently: how a device is found, whether the tuner is one frequency or a chain of elements, whether gain is a scalar or a set of named stages, what format samples arrive in, who owns the buffer. Application code inherits all of those answers. The demodulator and the decoder are portable and expensive to validate; the forty lines that open the radio are neither, and they are what make the rest device-specific.

Table 9 sets out the four ways a system can reach its hardware.

Table 9. Four ways to reach hardware, and what each costs. **The fourth row is the default for any system expected to outlive its first radio**, and the caveat in its second column is the subject of the rest of this section: an abstraction is only as good as what survived it.

Approach	What it gives	What it costs
Program the vendor interface directly	The deepest access to that vendor's hardware	One supplier for the life of the system
Wrap the vendor library yourself	Control over the abstraction, tailored to your needs	You now maintain an abstraction layer as well as an application
Use the framework's own source block	Works immediately inside that framework	Nothing outside that framework can use it: not a script, not a service, not a test rack
Use an open hardware abstraction layer	One interface, many devices, outside any single framework	A lowest common denominator by construction, and the omissions are the interesting part

THE FOUR THINGS AN ABSTRACTION LAYER WILL NOT CARRY

Absolute amplitude, because nothing in a generic interface says what a sample is worth in dBm. **Calibration**, because there is no standard call that returns a correction table or a calibration date, so a calibrated device and an uncalibrated one look identical through it. **Timing discipline**, because a generic call can report that a timestamp exists but never whether it is disciplined, or against what, so a free-running counter and a GNSS-locked one are the same field. **Device-specific measurement modes**, because a real-time density display or a pulse-parameter table has no generic representation. If your application depends on any of these, it depends on the hardware underneath rather than on the interface, and the abstraction has not made that dependency go away.

None of that is a defect. An abstraction that tried to standardize calibration across hardware that mostly has none would not have been adopted, and adoption is the whole value of an abstraction layer. But it does mean the question that decides whether an open framework can produce a defensible number is a question the abstraction cannot answer, and the buyer has to put it to the hardware instead. BNC-HB-001 documents one such interface in working detail, including the call that returns full scale and why it is the one the whole argument turns on.

26 The arithmetic that decides whether a system is buildable

Two calculations settle most architecture arguments before any hardware is chosen, and both are short enough to do in a meeting.

First, the data rate out of a receiver:

$$r = f_s \times 2 \times b/8 \qquad (7)$$

where

r is the data rate, in bytes per second

f_s is the complex sample rate

b is the width of the container each component travels in, which is not always the converter's resolution

At 125 MSPS with 16-bit components, equation (7) gives 500 Mbyte per second. That is more than a mechanical disk will take and about what a SATA solid-state drive will sustain, which is why a wideband capture architecture is a storage design as much as a radio design. An NVMe drive has several times that headroom, but it buys the headroom with power and heat, which is exactly the budget a battery-powered field instrument does not have spare. **The number that matters is the sustained rate, not the burst rate**, and section 16 explains why a device quoting only one figure has answered the easier question.

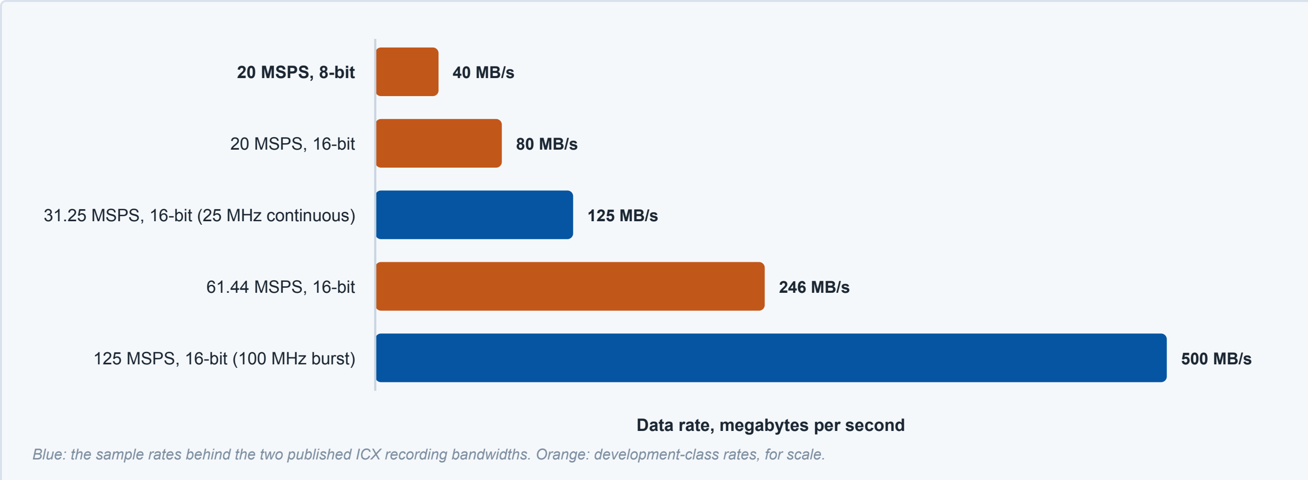

Figure 6. Equation (7) evaluated at rates this market actually runs. **Every bar is a sample rate, because that is the quantity the relation takes.** A published bandwidth has to be converted before it can be used here, by dividing by the 0.8 usable ratio of section 8: the blue bars are the 31.25 and 125 MSPS behind the published 25 MHz continuous and 100 MHz burst recording bandwidths. The first two bars share a sample rate and differ only in component width, which is the cleanest way to see the other term: halving the width halves the rate, so an 8-bit receiver is cheap to store as well as to buy. **Read these as hourly figures to feel the problem:** 100 Mbyte per second is 360 Gbyte an hour, and 500 Mbyte per second is 1.8 Tbyte an hour. An architecture that does not decimate at the source or trigger on something fills any disk you own, which is why the interesting engineering in a monitoring system is almost always in what it decides not to keep. DERIVED

Keep Figure 6 in view during any architecture discussion, because the storage question is usually raised last and decides the design.

Second, how long a buffer holds. A capture memory divided by the rate above gives the window you can take away and reprocess. For the ICX-FieldHawk's published 128 Mbyte of capture memory at 500 Mbyte per second that is about 0.26 seconds of gap-free 100 MHz capture, assuming a 16-bit component width, which is a conventional interleaved format rather than a published property of the instrument.

27 How a processing framework actually runs

The framework layer is where custom work lives, so it repays understanding one level below the block diagram. BNC-HB-001 documents this in depth; what follows is the part that changes a purchasing decision.

A graph-based framework connects processing blocks with buffered edges. Each block runs on its own thread, asks how many items are available and how much output space is free, and does work when both allow. Every output port gets one circular buffer, allocated when the graph starts and never resized. That single fact explains the behavior that surprises people most.

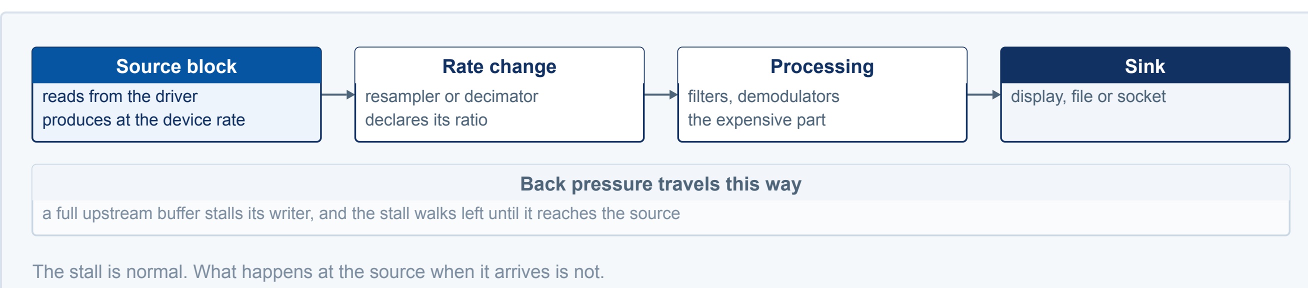

Figure 7. How a processing graph stalls, and why it matters for hardware selection. When any block cannot keep up, it stops consuming, the upstream block's output buffer fills, and that block waits for room rather than for data; the stall propagates left until it reaches the source. **With a file source, that is harmless and the graph simply runs slower. With a receiver, it is not**, because the converter keeps producing whether anyone is reading or not, so the driver's queue fills and samples are discarded. The stream downstream has no gap in it, because a gap would need the source to have produced something and it produced nothing. What the loss leaves is a discontinuity in time and phase that no downstream block can identify. SCHEMATIC

Figure 7 is the reason a device's transport and a host's performance are part of the measurement chain rather than logistics. The symptom is intermittent decode failure with a perfectly healthy-looking spectrum, and the cure is almost never a faster computer: it is decimating at the source so the samples are never produced, cutting display update rates, and taking per-sample work out of interpreted code.

> **WHAT TO ASK, GIVEN THAT**
>
> Whether the device can **decimate in hardware before the transport**, and by how much. A receiver that offers decimation from 1 to 4096 in powers of two lets you remove the problem at the source; one that streams at a fixed rate makes the host's ability to keep up part of your measurement. That is a specification question with a numeric answer, and it is not on most datasheets' front page.

28 Samples, scaling and time: the three that break integrations

An engineer told to integrate a receiver usually loses the first week to one of three things, and none of them appears in a datasheet. They are worth stating plainly because each has a symptom that looks like something else.

The format. Samples arrive as interleaved pairs, one in-phase value then one quadrature value, repeating. Two things about that sentence have to be confirmed rather than assumed: the order of the pair, and the type of each value. Sixteen-bit signed integers and thirty-two-bit floats are both common, and so is the byte order question, which does not arise until the day the capture host and the analysis host are different machines. **The symptom of a swapped pair is a spectrum mirrored about the center**, which is easy to mistake for an image problem in the radio.

> **THE SCALING TRAP, IN ONE NUMBER**
>
> A converter narrower than its container has to be positioned inside it. Twelve-bit samples placed in the top twelve bits of a sixteen-bit word read **sixteen times larger** than the same samples placed in the bottom twelve. That is 24 dB, and it is a clean, constant, entirely plausible-looking offset that no amount of staring at a constellation will reveal. Ask what full scale is, in the units the interface delivers, and check it once against a source of known level. **Full scale is a convention, not a fact**, and it is the single most common cause of an amplitude that is wrong by a suspiciously round number.

The time base. A timestamp is almost never attached to every sample; it is attached to the first sample of a buffer, and the time of every later sample is inferred by counting at the sample rate. That inference is what makes timing cheap, and it is also what makes a dropped sample dangerous. **A lost sample does not leave a hole. It shifts everything after it earlier by one sample period and leaves no evidence in the samples themselves.** The interface may still tell you, which is the whole point of asking whether it does. In a 125 MSPS stream one lost sample is 8 ns, which is nothing; a hundred thousand lost samples during one disk stall is 0.8 ms, which is more than enough to break a time-difference measurement or a correlation between two units.

> **WHY AN OVERRUN COUNTER MATTERS MORE THAN IT SOUNDS**
>
> This is the mechanism behind the gap-free argument of section 33, seen from the software side. **An interface that reports overruns lets you know the record is damaged; one that does not leaves you with a file that looks perfect and is wrong.** The report is usually a return code or a counter rather than anything visible, which is why section 52 tells you to read the stream return value rather than watch the console. Something is usually printed, since common source blocks emit a single character to standard error when a stream overruns, but a character on a terminal disappears with the terminal, is invisible to a headless service and never reaches the file. Note also that a negative return means the same thing, and **a quiet band is not one of them**: a receiver produces samples whether or not anything is transmitting, because noise is samples, so a read that times out has not found silence, it has found no data, which is a fault. An overrun is a different code again, so a wrapper that collapses every negative into one error hides the distinction you most need. Ask three questions of any interface before you build on it: does it report dropped samples, does it report them per buffer or only as a running total, and does the timestamp come from the hardware or from the host clock. A host-clock timestamp is a record of when software got around to looking, not of when the sample existed.

None of these three is difficult once named. All three are expensive when they are not, because each produces a plausible wrong answer rather than an obvious failure, and a plausible wrong answer can survive a long way into a project.

29 Platforms, and where RF data goes after the disk

Two practical questions close out the software part, and both are usually discovered after a purchase rather than before it.

What runs it. A receiver's usefulness is bounded by the machines its software supports. The ICX line publishes Windows 11, 10, 8 and 7, Debian 12, 11 and 10, and Ubuntu 24.04 back to 18.04, on x64 and AArch64. The AArch64 entry is the one worth noticing: it means a single-board computer can host the receiver, which is what makes an unattended node affordable, and it is the row most often missing from a competing datasheet. Ask for the list, and ask how long each entry is supported, because an operating system that falls out of support takes your deployment with it.

Where the data goes. Section 26 prices the disk; it stops there, and most systems do not. Three patterns are common and they have different failure modes. Keeping everything centrally is simplest and is bounded by the link, not the disk: a single sensor streaming continuously will exceed most site uplinks, so the architecture becomes a conversation about bandwidth rather than storage. Processing at the node and sending results is what makes a network of sensors affordable, and it moves the hard problem to deciding what a result is, because **whatever the node discards cannot be reanalyzed when the question changes**. Sending triggered captures is the middle path and the most common answer, which puts the weight back on section 34: the trigger becomes the compression algorithm, and its settings become part of the evidence.

> **THE QUESTION THAT DECIDES THE ARCHITECTURE**
>
> **Will anyone need to ask a different question of this data later?** If yes, something has to keep raw samples, and the cost of that is known in advance from the arithmetic in section 26. If no, process at the edge and keep results. Most projects answer no and discover a year later that they meant yes, which is expensive precisely because the samples are the one thing that cannot be reconstructed.

30 Diagnosing an overrun, and testing a pipeline

Two pieces of software practice are worth more than any amount of architecture advice, because they are what the first month actually consists of.

Triage an overrun in this order, because each step is cheaper than the one after it and rules out a whole class.

- **Confirm it is real.** Read the stream return value, not the console. A pipeline that reports nothing may be losing samples silently, which is the case section 28 warns about.
- **Ask whether it happens immediately or under load.** Immediately means a configuration error, usually a rate the host was never going to sustain. After minutes means a resource that fills or degrades: a disk buffer, a log, a thermal throttle.
- **Time the blocks individually.** One slow block stalls everything upstream of it through the back pressure of section 27, and the symptom appears at the source rather than at the culprit. Replace the suspect with a null sink and see whether the overruns stop.
- **Take the disk out of the loop.** Write to a null device at the same rate. If the overruns vanish, the problem is storage, and the fix is the arithmetic of section 26 rather than anything in the radio.
- **Then look at the host.** Power management that drops clock speed, an unpinned thread migrating between cores, memory that is not locked, and an interface with a shallow buffer are the usual remainder. **A machine that idles at low clock speed will overrun a stream it can easily sustain at full speed**, and that one accounts for a surprising share of pipelines that worked yesterday.

> **THE KNOBS BEHIND THAT LIST, SINCE A CONDITION YOU CANNOT ACT ON IS NOT ADVICE**
>
> Each step above has something specific to reach for, and the names are worth carrying even though the exact spelling belongs to your version. **Return codes:** an abstraction layer's read call returns a count when it succeeds and a negative constant when it does not, with distinct constants for an overflow and for a timeout; catch the constant, not the sign. **Buffer depth:** a framework block exposes a minimum output buffer size, and raising it on the source absorbs a scheduling hiccup that would otherwise become a dropped sample. **Clock speed:** set the host governor to its performance setting rather than its on-demand one, because a machine that idles slowly will overrun a stream it sustains easily at full speed. **Memory:** lock the process's pages so nothing in the sample path can be paged out. **Scheduling:** raise the priority of the thread that reads the device. **Full scale:** the abstraction layer exposes a call that returns the full-scale value for the current format, and using it rather than assuming the container width is what avoids the 24 dB error above.

> **HOW TO TEST A RECEIVING PIPELINE WITHOUT HARDWARE ON EVERY DESK**
>
> **Record once, test forever.** A capture with known contents, kept in the repository, turns an RF pipeline into ordinary software: the same file in, the same numbers out, on every commit. Ask a vendor for one as a deliverable rather than a favor. Two tests are worth building on it. **An amplitude test**, where a capture of a known level must come out at that level within a stated bound, which catches the scaling error of section 28 the day it appears rather than the week before a report. And **a gap test**, where a capture with a known sample count must produce exactly that many samples, which catches a silent drop that no spectrum display would show.

One representation detail belongs here because it is invisible until it is expensive. **A timestamp carried as a floating-point count of seconds since the epoch cannot hold nanoseconds.** A double has about sixteen significant digits, and at a present-day epoch value that leaves about 238 nanoseconds of resolution, which is thirty samples at 125 MSPS. For a spectrum display that is irrelevant. For time-difference work between two units it is the entire measurement. Carry integer seconds with an integer fraction, or a 64-bit count of nanoseconds, which spans a couple of centuries and loses nothing.

31 What a software lead should ask before anyone commits

The person who will carry the integration is usually not in the room when the receiver is chosen, and the questions that decide their year are not on any datasheet. Six are worth asking early enough to matter.

- **What language is the fast path in, and where does the interpreter sit?** A graph assembled in Python whose blocks are compiled is fine, because the interpreter only builds the graph. A custom block written in Python sits in the sample path and takes the interpreter's global lock with it, so it serializes against every other Python block in the process. This is the most common reason a pipeline that worked on a bench recording will not keep up with a live receiver.
- **How many times is a sample copied between the converter and the first useful operation?** Driver to library, and library into the framework's buffer, are usually copies. What happens after that often is not: a well-built framework hands a block a pointer into its upstream's output buffer rather than a copy, which is why the buffer model in section 27 is worth understanding. The copies that remain are the ones you pay for, and at 500 Mbyte per second memory bandwidth is a real budget rather than a rounding error.
- **What is the driver's dependency and version story?** An abstraction layer loads modules built against a particular version of itself, so a distribution upgrade can remove a working radio from a working machine without anyone touching the radio. Ask which versions are supported, how long, and whether a build from source is a supported path or a rescue operation.
- **How do we test this without a signal generator on every desk?** A pipeline that can be fed from a recorded file is testable in continuous integration; one that can only be fed from hardware is testable when somebody is in the building. Ask for a recorded capture with known contents, and treat it as a deliverable rather than a favor.
- **What is the support window, and how are updates delivered and signed?** The Cyber Resilience Act of section 54 turns this from a courtesy into a regulated commitment, and it is the question that decides whether the device is still deployable in five years.
- **What is the calibration, concretely, on disk?** If correction is applied before the samples reach you, then it is a file that was installed alongside the driver, it is specific to that unit, and it has a date on it. Ask where it lives, what happens if it is missing, whether the software warns or silently continues, and what the update path is after a recalibration. **A correction that can go missing quietly is a correction you cannot rely on**, and this question is rarely asked until the first time it happens.

> **WHY THESE BELONG IN THE PURCHASE CONVERSATION**
>
> Every one of them is cheap to answer before a decision and expensive to discover after one. None of them is an RF question, which is exactly why they get left out of requirements written by RF engineers, and why the vendor form in section 63 asks for a software contact as well as an applications engineer.

PART IV

From samples to measurements

The boundary that divides this market, what it costs to cross, and where the ICX-FieldHawk sits relative to it.

32 What separates a reading from a measurement

A number without a stated uncertainty and a traceable reference is a reading. A number with both is a measurement. Everything in Parts I to III is about producing readings well. This part is about the other thing, and the distinction is the single most consequential one in the handbook.

An uncorrected receiver reports its converter's output as a fraction of full scale. Turning that into an absolute power needs a correction table covering every gain state and every frequency in use, built against a reference source, maintained as it drifts and re-checked when anything in the chain changes. Nothing prevents an engineer from building that table. It is a large, dull, recurring job that most projects discover late.

> **CALIBRATION IS A PROPERTY OF THE PATH, NOT OF THE PROCESSING**
>
> It has to be applied to the samples **before** any processing block sees them. A correction applied downstream of a decision that was made on uncorrected data does not recover the decision. That is why the question is not whether a vendor supplies a correction table, but whether the correction is already in the samples when they cross into your code.

This is also why the absence of an amplitude accuracy figure on a datasheet matters more than it looks. A receiver quoting plus or minus 2.0 dB is making a claim about every reading it produces under stated conditions. A board quoting nothing is not making a weaker claim; it is making no claim, and the difference between those two positions is the difference between a result that can be published and one that cannot.

33 Real-time capture, and what gap-free guarantees

The second dividing line is what happens to signals that are not on all the time. A swept receiver visits each frequency in turn, and anything happening elsewhere while it looks here is absent from the result. The absence is silent: the trace looks the same whether the band was quiet or the receiver was pointed the other way.

A real-time receiver holds a span continuously and transforms every sample. Where the governing relations are published rather than only their results, a reader can price any configuration:

$$\mathrm{POI} = 2 \times N \times D \times 8\ \mathrm{ns} \qquad (8)$$

where

POI is the observation length that guarantees full-amplitude capture, in seconds

N is transform size, in points

D is decimation factor

$8\ ns$ is the engine's sample interval, the reciprocal of its 125 MSPS sample rate

The factor of two in equation (8) is not margin. A burst exactly one frame long, arriving halfway through a frame, puts half its energy in each of two frames and is measured at full amplitude in neither, so guaranteeing full amplitude means guaranteeing that a whole frame falls inside the burst.

That coefficient belongs to contiguous, non-overlapping frames, which is the arrangement a relation of this form describes. An engine that overlaps its transforms needs less: where a new frame begins every half frame, any burst lasting one and a half frame lengths already contains a whole one. So the leading coefficient is part of the relation a vendor has to publish rather than a constant of nature, and a competitor quoting 1.5 is not necessarily flattering the number. Asking which coefficient applies is the second question, after asking for the relation at all.

Table 10 evaluates it at five settings of the ICX-FieldHawk real-time engine.

Table 10. The trade, evaluated. Each row is an observation length that guarantees full-amplitude capture, against the transform bin spacing at the same setting. **The last column is the point: multiply the two and the answer is exactly 2 every time**, because intercept goes as *N D* and bin spacing as its reciprocal, so no setting improves both. It is a dimensionless constant, and you can check any row of it on a calculator. Read the table as a purchase criterion: pick the shortest event the system must measure honestly, and that fixes the largest transform the configuration may use. **One condition the columns do not carry: decimation divides the span as well as the rate.** The first four rows all cover the full 125 MHz the engine digitizes, because *N* times the bin spacing is 125 MHz in each. The last row does not: at *D* = 64 the transform covers 1.953 MHz, so its guarantee is a guarantee over that window and not over the whole analysis bandwidth. **The first and third rows are the two published operating points**; the rest are the relation evaluated at other settings, and are marked as derived for that reason.

Transform *N*	Decimation *D*	Bin spacing	Guaranteed intercept	Product
32	1	3.90625 MHz	0.512 µs	2
256	1	488.28125 kHz	4.096 µs	2
2048	1	61.03516 kHz	32.768 µs	2
4096	1	30.51758 kHz	65.536 µs	2
2048	64	953.674 Hz	2.097152 ms	2

The last row is the end of the published range rather than an arbitrary stopping point, and the reasoning shows how to check a datasheet against itself. The published real-time resolution bandwidth runs 7.81 MHz down to 1.90 kHz in 13 grades under this window. The grid behind those grades is the thirteen powers of two in *N D* from 32 to 131,072, whose bin spacings at this engine are 3.90625 MHz down to 953.674 Hz. **The column that confirms the grid is the flat-top one published beside it**: 14.73 MHz divided by 3.90625 MHz is 3.771, and 3.59 kHz divided by 953.674 Hz is 3.764, both within 0.2 percent of the flat-top window's equivalent noise bandwidth of about 3.77. So **2.097 ms is the longest guaranteed intercept the published range reaches, over the 1.95 MHz span that decimation leaves**, and a requirement written past it is a requirement for a different instrument. Write both numbers into any specification that uses it, because an intercept guarantee without the span it applies over is the same half-answer this handbook keeps warning about.

ONE LOOSE END, REPORTED RATHER THAN SMOOTHED OVER

The Blackman-Nuttall column does not close as tightly. Its published endpoints sit about 1.2 percent above 1.9761 times the bin spacing, at almost exactly twice it, while the real-time display reports 1.9761 times the bin spacing directly, as section 13 shows. A 1.2 percent discrepancy in a three-significant-figure specification row is far more likely to be rounding in the datasheet than a different grid, and the flat-top column agreeing to 0.2 percent says the grid is right. **We are flagging it rather than quietly picking the number that suits the argument**, and it is listed in the verification note.

The first row is the purchase criterion at its sharpest. At 0.512 microseconds of guaranteed intercept, every event in the band lasting longer than about half a microsecond is measured at its true amplitude rather than glimpsed. A swept receiver at the same span offers no such boundary at all, and BNC-AN-102 works through how far apart the two answers land for a monitoring system.

34 Triggering, and how a finite buffer sees a long day

The arithmetic of section 26 says a wideband receiver fills any disk you own. The way out is not a bigger disk; it is deciding what to keep, and that decision is triggering. It is the least specified capability in this market and one of the most consequential.

Table 11 is the set of modes worth knowing by name.

Table 11. Trigger modes. **Ask two numbers about any of them:** how much of the record is kept from before the trigger fired, and how long the instrument is blind after one capture before it can arm again. The first decides whether you see what caused the event; the second decides what you miss while writing the last one away.

Mode	What starts the capture	What it is for
Free run	Nothing; samples stream continuously	Streaming to a process that decides for itself. The default, and the one that fills disks
Level	The signal crossing an amplitude threshold	Pulsed and bursty emitters. Simple, and prone to firing on noise if the threshold is set from a quiet band

Mode	What starts the capture	What it is for
External	A hardware edge on an input line	Synchronizing to something outside the receiver, and the only mode that ties two instruments to the same instant
Frequency mask	Energy entering or leaving a drawn region of the spectrum	**The one worth asking about.** It fires on a signal appearing where nothing should be, which is the definition of the event most monitoring jobs exist to catch, and it can ignore a strong legitimate carrier sitting alongside

> **PRE-TRIGGER DEPTH IS THE SPECIFICATION NOBODY ASKS FOR**
>
> An event that fires a trigger has usually already started. **If the instrument keeps only what follows the trigger, the interesting part is the part you threw away**, which is the rise, the turn-on transient and whatever preceded the emission. A capture memory running continuously as a circular buffer can hand you the milliseconds before the trigger; one that starts writing when the trigger fires cannot. This is a property of the acquisition architecture and it is rarely on a datasheet, so it belongs on the list in section 63.

35 Distributed systems, and what changes when there is more than one

A single receiver has one clock, one reference and one place. A network of them has as many of each as it has nodes, and almost every hard problem in distributed sensing comes from that multiplication rather than from the radios.

Table 12 sets out what changes.

Table 12. What multiplies when a system has more than one receiver. **The last row is the one that costs the most**, because a single instrument that misbehaves is noticed and a node that misbehaves at the far end of a site is not. Design the network so that each node reports enough about its own state, reference lock, temperature and overflow count, to be doubted from the center.

Quantity	One receiver	Many receivers
Frequency reference	Accuracy against its own oscillator	Every node must agree, or measurements cannot be compared across the network
Time	Relative within one capture	Absolute and common, or a time-difference calculation is meaningless
Amplitude	One bound	The same bound at every node, or a coverage map is a map of calibration errors
Geometry	Not applicable	Dominates. Dilution of precision multiplies whatever timing error survives
Transport format	A convenience	A commitment. Changing it later means changing every node
Failure	Visible. The instrument is in front of you	Silent. A node with a drifted reference keeps reporting plausible numbers

> **THE ARITHMETIC THAT DECIDES A DISTRIBUTED DESIGN**
>
> Position error is the geometric dilution of precision multiplied by the speed of light multiplied by the arrival-time uncertainty. At 0.2998 m per nanosecond, a hundred nanoseconds is 30.0 m of range uncertainty before geometry, and a workable rather than good layout at a dilution of three makes that about 90 m of position error. **Geometry contains no receiver parameter at all**, which is why a site survey moves this number further than any receiver purchase does. BNC-AN-110 works the full budget.

36 How a swept analyzer works, and the relation that governs it

The antagonist of this part deserves a fair description, because it is the right answer to a great many problems and because its limitation follows from a relation as clean as the one on the other side.

A swept-tuned analyzer moves a narrow filter across the span and records what comes through it. It sees one resolution bandwidth at a time, and the filter needs time to settle: a filter of bandwidth RBW cannot deliver a settled output faster than about 1/RBW. Dwelling that long on each of Span/RBW positions gives the sweep time.

$$T_{\text{sweep}} \approx k \times \text{Span} / \text{RBW}^2 \quad (9)$$

where

T_{sweep} is the time to cross the span once, in seconds

$Span$ is the frequency range swept, in hertz

RBW is the resolution bandwidth, in hertz

k is a constant near 2 set by the filter shape

> THE INVERSE SQUARE IS THE WHOLE STORY
>
> **Equation (9) is quadratic in resolution bandwidth**, so the setting that finds weak signals is the setting that makes the sweep slow. Ask for ten times finer resolution and the sweep takes a hundred times longer. A 100 MHz span at 1 kHz resolution is around 200 seconds; the same span at 30 kHz is a fifth of a second. **That is the trade a swept-filter architecture cannot escape**, and it is why the honest comparison is not real-time against swept but 'what fraction of the time is the instrument looking where the event is'.

One qualification, because it is the objection this argument attracts. The relation above is the analog swept-filter law. An analyzer that steps a transform across the span instead needs a record length of about 1/RBW per segment, with the segment count set by its analysis bandwidth rather than by its resolution bandwidth, so its sweep time falls to roughly Span / (B × RBW), where B is the analysis bandwidth of each segment: linear in 1/RBW rather than quadratic, and often far faster than the swept-filter law predicts. The contrast is not subtle. A 40 GHz span at 1 kHz resolution is about 80,000 seconds under the swept-filter law and about 0.4 seconds for a stepped transform with a 100 MHz analysis bandwidth. **What no stepped architecture escapes is the second half.** While it is looking at one segment it is not looking at the others, so the argument about what fraction of the time the instrument is pointed where the event is survives intact even where the speed argument does not.

Set that beside the real-time relation and the two architectures separate cleanly. A swept instrument's cost goes as the square of the resolution you ask for in a swept-filter design and linearly in a stepped-transform one, and either way it is blind to everything outside its filter while it sweeps. A real-time engine holds the whole span continuously, and its cost is a fixed intercept guarantee that goes linearly in transform size. **Neither is better; they are priced in different currencies.** A swept analyzer is the right instrument for a stable signal you can dwell on, for the best available resolution and dynamic range at a price, and for spans wider than any real-time engine digitizes at once. It is the wrong instrument for anything that is not there when you arrive.

37 The evidence a swept trace cannot give you

Section 33 argued the point from a relation. This section argues it from a screen, because the difference between a receiver that holds a band and one that visits it is visible without reading a single number.

Figure 8. The 2.4 GHz band held whole: 2.39 to 2.49 GHz in one 100 MHz window at 30 kHz resolution, with max hold and a single sweep overlaid in the center pane, a spectrogram above and a 31.67 MHz zoom below. **Read the gap between the two traces.** The upper trace is what accumulation finds; the lower one is one instant. Where they differ, by 20 dB and more in places, is exactly what a receiver that visits this band rather than holding it would have reported as empty, and the difference is a function of when it happened to look rather than of what was there. **MEASURED**

Figure 8 is the argument of Part IV in one picture, and it generalizes past this band. Any environment with intermittent or hopping emitters produces the same divergence, and the size of it is set by duty cycle rather than by anything in the receiver's specification. **A single-sweep trace of a bursty environment is not a measurement of that environment; it is a sample of it at an instant nobody chose.** The spectrogram above the traces is the same information in time, which is often the faster way to see that an emitter is present at low duty cycle rather than absent.

> **WHAT TO ASK A VENDOR ABOUT THIS**
>
> Not whether the device has a real-time mode, which almost everything now claims, but what its **guaranteed intercept** is and at what transform size. If the answer is a marketing phrase rather than a relation you can evaluate, the honest reading is that the guarantee does not exist. Equation (8) is what a published relation looks like: it lets a buyer price a configuration the vendor never considered.

38 What the open world does with a calibrated front end

The two halves of this handbook meet here. Part III described an open software ecosystem that can decode almost anything; Part IV described a class of hardware that produces numbers you can defend. The interesting question is what happens when they are put together, and it has been demonstrated end to end rather than argued.

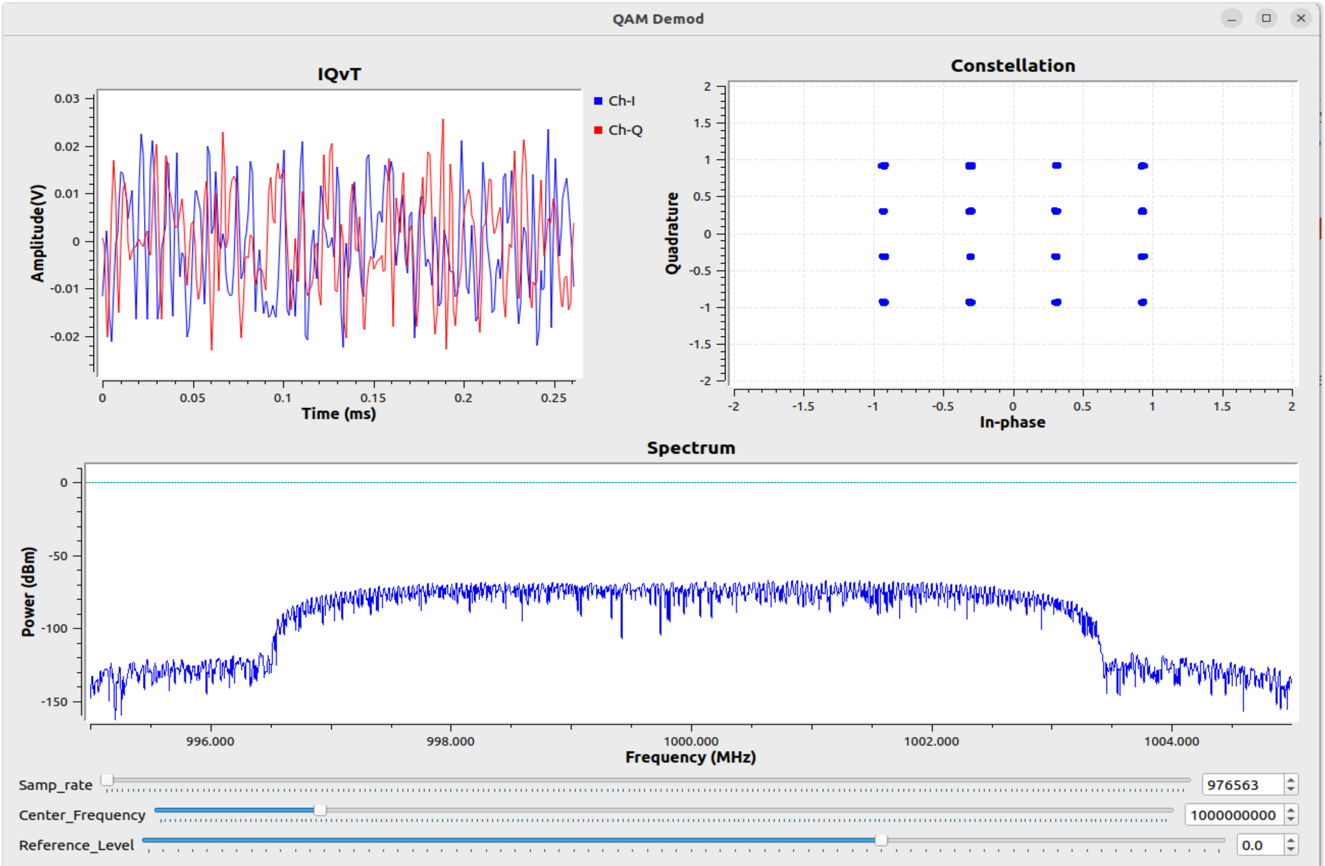

Figure 9. A 16-QAM constellation recovered in an open signal-processing framework: sixteen clusters on a four-by-four lattice at three distinct amplitude levels, with the IQ record and the received spectrum alongside. The controls read a 1 GHz center, 976,563 SPS and a 0 dBm reference level; the −80 dBm source level quoted below is the integration guide's figure and is not legible on this display. **One trap on the spectrum pane, and it is the kind this handbook trains you to catch.** Its axis runs about 995 to 1005 MHz, roughly 10 MHz, while a complex rate of 976,563 samples per second represents 976.563 kHz. The axis is the display block's own setting, not the acquisition span, so read it as a plot control rather than as evidence of bandwidth. The rate itself is exact: 976,562.5 is 125 MSPS divided by 128, displayed rounded. **Read what the calibration does and does not buy here.** Cluster tightness is error vector magnitude, set by noise figure, phase noise and loop bandwidths, and the calibration does not improve it. What the calibration buys is the axis the constellation is drawn on, and an absolute level for the carrier that produced it. An uncorrected board draws the same picture and cannot tell you what −80 dBm means. MEASURED

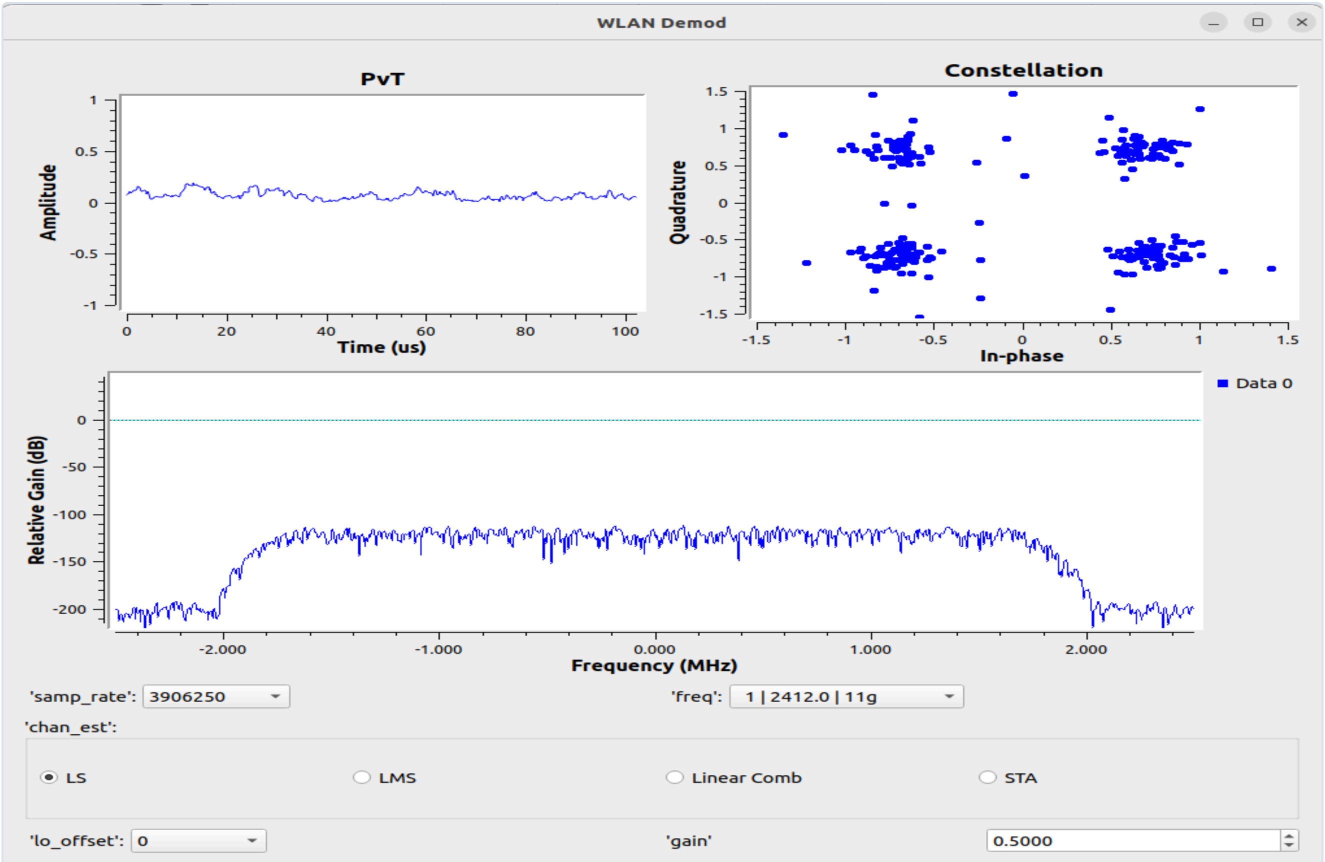

Figure 10. The 802.11a orthogonal frequency-division multiplexing physical layer, which carries data on many narrow subcarriers at once, recovered at 2.412 GHz in the same framework, with the channel selector reading 11g and a least-squares channel estimator selected: four clusters in the constellation pane and the received spectrum below. Read this window with the same caution as the last one. The sample-rate field reads 3,906,250, an exact binary decimation of the converter rate, but the spectrum axis runs to about ±2.5 MHz and the signal fills roughly 4 MHz of it, both wider than that rate represents. The axis is again a display setting. It follows that the width the signal appears to occupy on it cannot be read as a bandwidth either, and the source material does not state which channel configuration produced the capture. **So no bandwidth claim is made from this window**, and none is needed: what it shows is that the frames decoded, and they leave as a packet capture file rather than a picture. **MEASURED**

Figure 9 and Figure 10 are the argument of this handbook in two screenshots. Both were produced by open blocks that were never written for this instrument, attached at a standard boundary, with the amplitude correction already applied beneath them, a behavior demonstrated in the integration guide these captures come from rather than specified on a datasheet. **Neither half gave anything up to work with the other**, and that is the unusual position: most of this market is a choice between open and programmable hardware that is not calibrated, and calibrated instruments that will not let you inside.

> **THE PURCHASING QUESTION THIS SECTION IS REALLY ABOUT**
>
> Not whether a device works with open software, which almost everything now claims, but **what crosses the boundary with the samples**. If the answer is a fraction of full scale, you have bought a receiver and the amplitude work is ahead of you. If it is a level in dBm with a bound behind it, you have bought a measurement path, and the difference is worth exactly what the correction table would have cost you to build and maintain.

39 Calibration, and what traceability actually costs someone

Part IV keeps returning to calibration, so here is what the word covers, because it is used loosely enough to have almost stopped discriminating.

- **A correction** is a number applied to a reading. Anyone can produce one, and a correction with no provenance improves precision and does nothing for accuracy.
- **A calibration** is a correction established by comparison against a reference of known accuracy, at a known date, under known conditions. It has an owner and an expiry.
- **Traceability** is the unbroken chain from that reference to a national or international standard, with a stated uncertainty at every step. It is what lets a number be

defended by somebody who was not present when it was taken.

The chain is the expensive part, and it is why calibration has resisted commoditization while everything around it has become software. Somebody has to own reference equipment, have it calibrated against something better on a schedule, maintain an environment in which the comparison means something, and keep records that survive an audit. **None of that scales down**, which is precisely why a device class that does not carry it exists and is the right answer for most jobs.

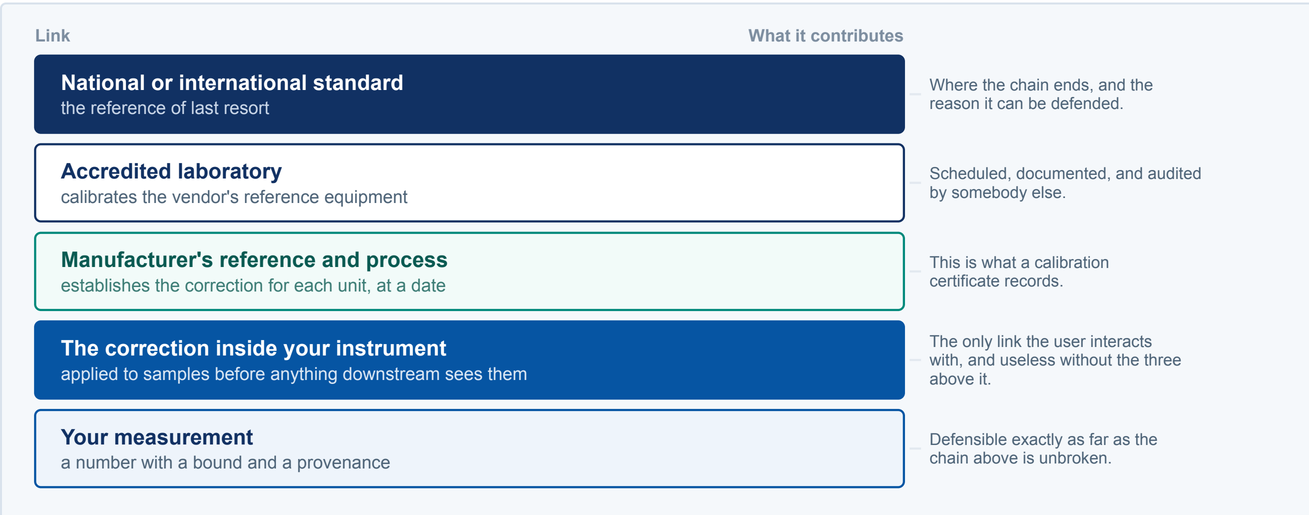

Figure 11. The traceability chain, top to bottom. **Every link has to hold**, and the user interacts only with the fourth. This is why an amplitude accuracy figure on a datasheet is a claim about the three links above it rather than about the instrument alone, and why a device that publishes no such figure is not making a weaker version of the same claim. It is making no claim, and no amount of processing downstream reconstructs one. SCHEMATIC

Figure 11 also answers a question buyers ask in good faith: why not calibrate a cheap receiver yourself? You can, and for a single fixed configuration it is a reasonable afternoon's work. What does not scale is the second unit, the second frequency, the second gain state, and the drift that starts the day after. The correction table is not the hard part. **Maintaining it, across every state the device can be in, for as long as the program runs, is the hard part**, and that recurring cost is what a calibrated instrument is actually selling.

40 Machine learning on spectrum, and what it inherits

A growing share of what people want from a wideband receiver is not a display but a training set. Modulation recognition, anomaly detection over a monitored band, and radio frequency fingerprinting, which identifies an individual transmitter from the small imperfections of its hardware, are all now ordinary asks. The handbook has already supplied everything needed to see where they go wrong.

> **THE IMPAIRMENTS OF SECTION 12 ARE THE FINGERPRINT**
>
> Radio frequency fingerprinting works precisely because a transmitter's quadrature imbalance, carrier leakage, phase noise and amplifier nonlinearity are individual to the unit. **Those are the same quantities the receiver also has.** A model trained on captures from one receiver can learn the receiver's impairments instead of the transmitter's, and it will score well in testing and fail on the day the hardware is swapped. Published work on this has been arguing about exactly that failure for a decade, and it remains open.

This is the handbook's argument arriving in a new place. A model is a measurement instrument built out of data, and it inherits every property of the data it was built from. If the amplitude axis of a training capture has no absolute meaning, the model has learned relative structure and cannot transfer to a receiver whose gain differs. If the captures are not gap-free, the model has learned from a record with holes in it that are invisible in the file. If the time base drifts, a model that uses timing features has learned the drift.

Four practices follow, and none of them is about the model.

- **Record raw samples, not processed traces.** A spectrogram image discards phase and throws away most of what a fingerprinting model uses. Keep the samples and the metadata that makes them interpretable.
- **Train across receivers, not just across transmitters.** If every capture comes from one unit, the model cannot distinguish the transmitter from the receiver, and no

amount of held-out test data will reveal it, because the held-out data has the same receiver in it.

- **Calibrate before you collect, not after you fail.** A correction applied to the samples makes the amplitude axis comparable across units, days and temperatures, which is what makes a dataset reusable rather than disposable.
- **Keep the provenance with the data, in a format somebody else can read.** Receiver, serial number, gain settings, sample rate, antenna, date, temperature, calibration due date. The field has a convention for this, the Signal Metadata Format, which pairs a binary sample file with a JSON sidecar describing it; using it costs nothing and is the difference between a dataset and a directory. A capture without provenance is not a measurement and it is not a training example either.

None of this argues for or against any particular instrument. It argues that the dataset is the deliverable, and the properties that make a measurement defensible are the same ones that make a dataset transfer.

41 Where the ICX-FieldHawk sits

The ICX-FieldHawk is a real-time spectrum analyzer family built around one analysis engine and one host interface, SpectraCore, in three physical forms. It belongs in the fourth row of Figure 4, and this section says what that means in published numbers rather than in adjectives.

Table 13 is the family.

Table 13. The published family. Two things read differently against Part I. The frequency column reaches well past the 6 GHz ceiling that the integrated-transceiver architecture imposes, which means a different front end and is most of the reason this is a different class of product. And **only the 9.5 GHz and 40 GHz tiers carry published specification tables**, so any requirement written against a number in this handbook has to fall back to one of those two models.

Model	Form factor	Frequency	Analysis bandwidth
ICX-400	Handheld, 10.1 in multi-touch display	9 kHz to 40 GHz	100 MHz
ICX-090R	IP68 rugged tablet	9 kHz to 9.5 GHz	100 MHz
ICX-400R	IP68 rugged tablet	9 kHz to 40 GHz	100 MHz
ICX-090U	USB or LAN module	9 kHz to 9.5 GHz	50 MHz std, 100 MHz optional
ICX-400U	USB or LAN module	9 kHz to 40 GHz	100 MHz

Source: the ICX-FieldHawk handheld, rugged and USB datasheets.

Table 14 is the performance, each row with the condition that makes it true.

Table 14. Published performance with conditions. Two rows answer questions Part II said most SDR datasheets leave open. **Amplitude accuracy is stated as a bound**, which is the field whose absence defines the categories above it. And **recording bandwidth is published as two numbers** rather than the flattering one twice, which is the behavior section 16 told you to look for.

Parameter	Published value	Condition
Amplitude accuracy	±2.0 dB / ±3.0 dB	9 kHz to 9.5 GHz / 9.5 to 40 GHz
Displayed average noise level	−159.9 dBm/Hz (400); −167.5 dBm/Hz (090)	1 GHz, RBW 1 kHz
SSB phase noise	−107.5 dBc/Hz (400); −101.6 dBc/Hz (090)	1 GHz carrier, 10 kHz offset
SSB phase noise	−85.7 dBc/Hz	40 GHz carrier, 10 kHz offset
Third-order intercept	+40.3 dBm	1 GHz, reference level +20 dBm, ICX-400
Real-time engine	frame rate = $10^9 / (N\,D \times 8)$ frames/s	FPGA transform, no missing samples
IQ capture	125 MSPS, decimation 1 to 4096 in powers of two	128 Mbyte capture memory
Recording bandwidth	100 MHz burst; 25 MHz continuous	published as two figures
Frequency reference	TCXO < 1 ppm; OCXO < 0.15 ppm	OCXO is option 01
GNSS 1PPS	±100 ns standard	±75 ns / ±50 ns with options

Source: the handheld, rugged and USB datasheets. Standing conditions: 10 minute warm-up, 25 °C, spur reject standard.

42 What an instrument does that a receiver does not

The abstract difference is calibration. The concrete difference is that an instrument carries measurement functions whose results are defined, repeatable and checkable from the screen. That last property is worth more than it sounds, and it is the practical test to apply to any device claiming instrument status.

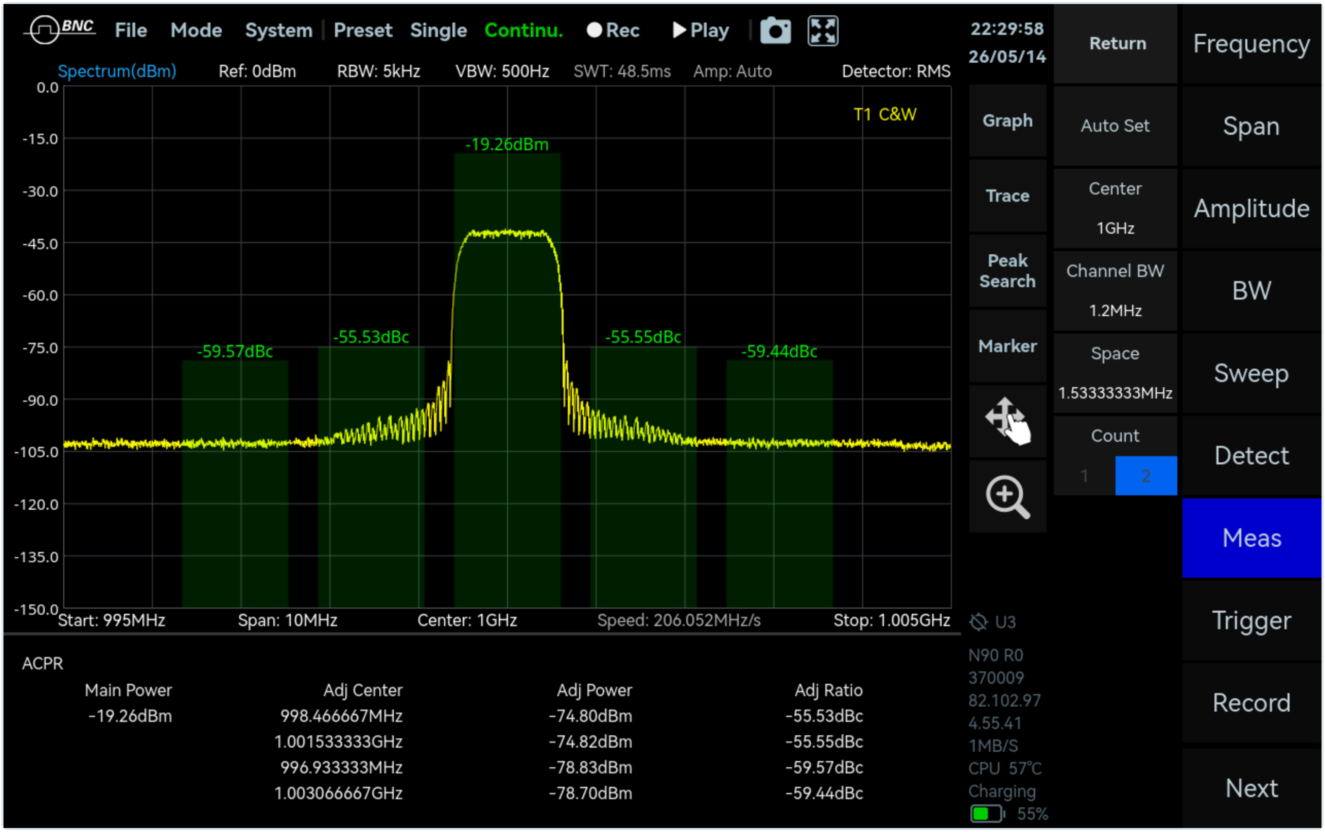

Figure 12. Adjacent channel power ratio on a 1 GHz carrier: main channel −19.26 dBm in a 1.2 MHz channel bandwidth, adjacent channels at ±1.533333 MHz spacing, and four ratios reported from −55.53 to −59.57 dBc. **Every number here is recoverable from the others.** Subtract the main power from each adjacent power and you get the ratio column to within 0.01 dB, which is the last digit each field is independently rounded to, so two of the four land one count away rather than exactly; add and subtract one and two channel spacings from 1 GHz and you get the four adjacent centers exactly. Settings that would be easy to leave wrong are on the glass too: 5 kHz resolution bandwidth, 500 Hz video bandwidth, and an RMS detector, which is the combination that measures a modulated signal's power rather than its peaks. MEASURED

Figure 12 is the property to look for, stated generally: **a result you cannot recompute is a result you cannot defend.** A device that reports a derived quantity and also reports the raw quantities it derived it from lets a reader check the derivation without trusting the device. That is a design choice rather than an accident, and it is rare below the instrument class because nothing else in the stack is obliged to make it.

Those functions run under SpectraCore, the host software shared across every form factor in the family, and the published list behind that behavior is long and, more usefully, is included rather than licensed piece by piece: channel power, occupied bandwidth, x dB bandwidth, adjacent channel power ratio, third-order intermodulation, spectrum emission mask, AM and FM demodulation, antenna factor, amplitude offset, signal track, peak table, data record and playback, multiple unit display and amplitude correction. The operating modes are standard spectrum analysis, IQ streaming with digital down-conversion, power detection at 8 ns time resolution, real-time spectrum, phase noise measurement from 1 Hz to 10 MHz offsets, and harmonics analysis. Digital demodulation and pulse detection are options 71 and 72.

> **THE COMMERCIAL ARGUMENT, MADE FROM THE PUBLISHED LIST ONLY**
>
> A fleet assembled from these units is uniform, because a routine written against one runs against all of them: the measurement set is not licensed function by function. That argument is made here from what is published and never from a claim about anyone else's price list, which is a claim this handbook is not in a position to verify and would not print if it were.

43 Four more measurements, each checkable from its own screen

The ACPR screen made a general point. These four make it four more times, and together they are the practical answer to what an instrument buys over a receiver: not prettier traces, but derived quantities you can verify without trusting the instrument that produced them.

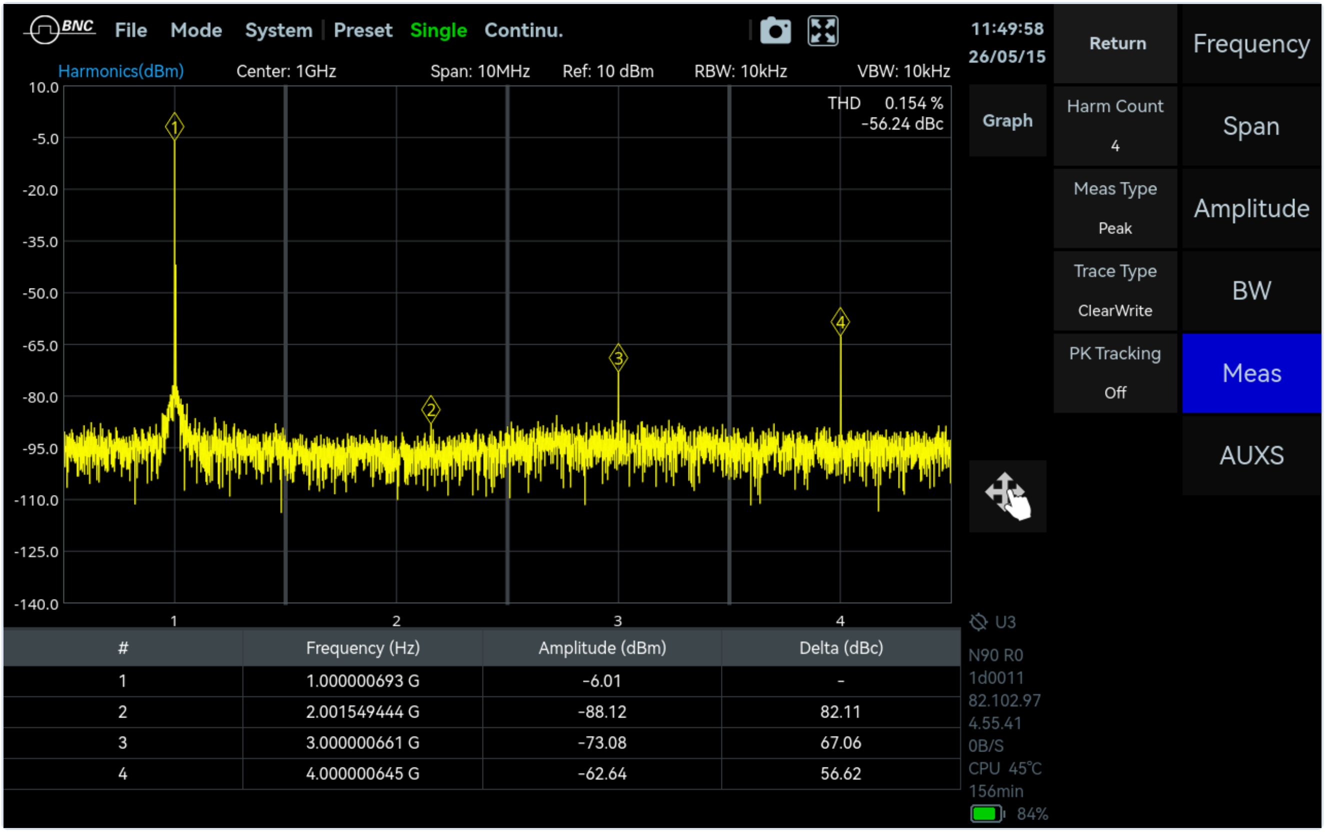

Figure 13. Harmonics analysis of a 1 GHz source in a 10 MHz span at 10 kHz resolution and video bandwidth, clear-write trace, peak measurement, reference level 10 dBm: fundamental at −6.01 dBm, second harmonic 82.11 dB down, third 67.06 dB down, fourth 56.62 dB down, with total harmonic distortion reported as 0.154 percent and −56.24 dBc. **Both readings are the same number and both come out of the table above them:** take the amplitude column, subtract the fundamental from each harmonic, root the summed power ratios and you get 0.1541 percent, whose log times twenty is −56.243 dBc. Work instead from the delta column exactly as printed and you get 0.1543 percent and −56.23 dBc; the hundredth of a decibel between the two routes is the rounding in the columns, not a disagreement, and it is worth knowing which route you took before reporting the digit. Two details reward a second look, and the second is the more useful. **The harmonics rise with order rather than decaying**, the fourth sitting 10.4 dB above the third, so a single distortion figure would have pointed at the wrong corrective action. **And the second-harmonic marker is not a harmonic.** It reads 2.001549444 GHz against a fundamental of 1.000000693 GHz, so it sits 1.548 MHz off exactly twice the fundamental, and at −88.12 dBm it is about seven decibels above the grass: that is a peak search finding noise. Markers three and four land within 2 kHz of exact. The total survives it, because that term carries under a third of one percent of the summed power, but **a reader who checks the frequency column finds this before we do, and a caption that did not say so would be the weaker document.** MEASURED

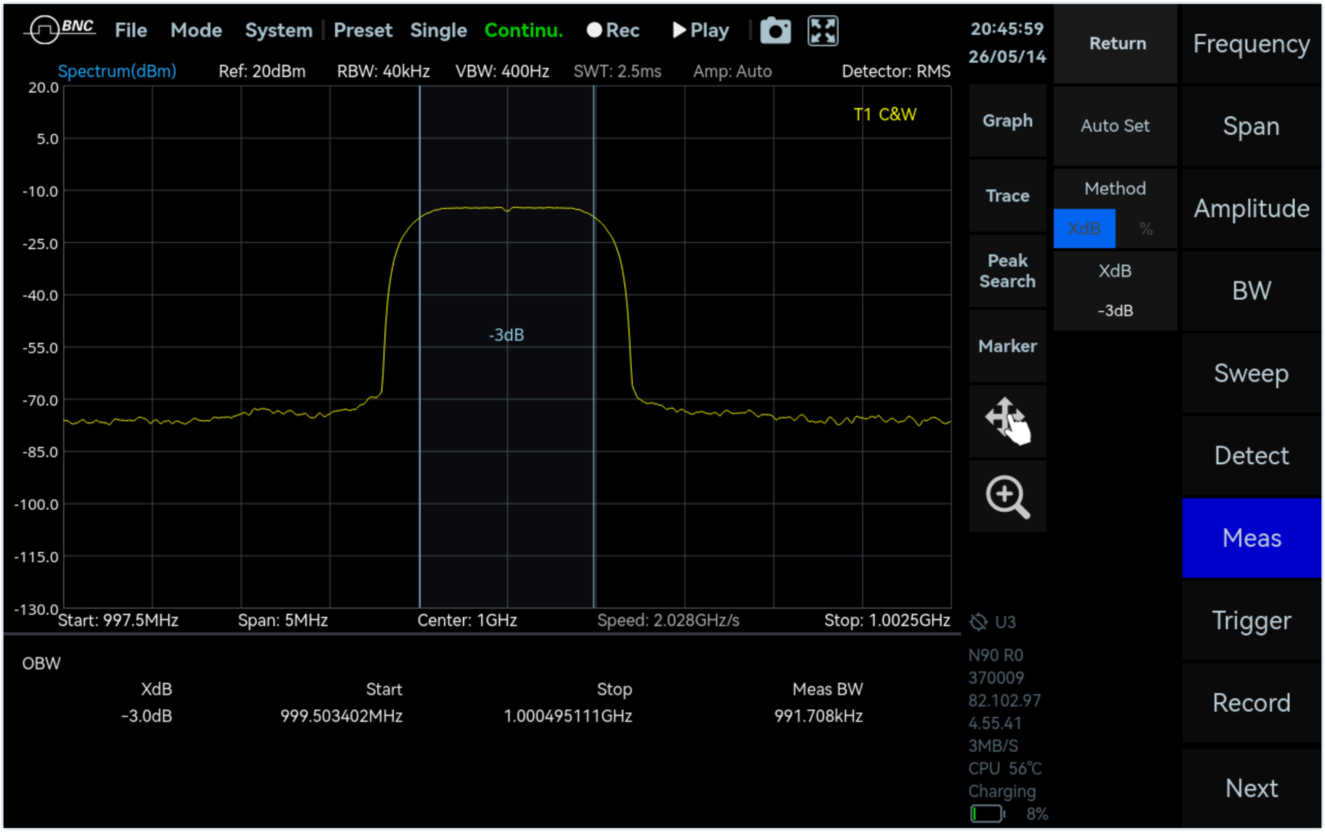

Figure 14. Occupied bandwidth by the x dB method at −3.0 dB on a shaped 1 GHz carrier: band edges reported at 999.503402 MHz and 1.000495111 GHz, measured bandwidth 991.708 kHz. Subtract the edges and you get 991.709 kHz, one hertz apart in nine hundred thousand, which is the rounding in the edge readings rather than a disagreement. The settings that make the result mean anything are on the glass too: 40 kHz resolution bandwidth, 400 Hz video bandwidth, and an RMS detector again. **One caution on the word.** This is the instrument's x dB function, and what it reports is a −3 dB width. Regulators generally use occupied bandwidth to mean the band containing 99 percent of the total power, which is a different and usually wider number on the same signal. Both are legitimate; quoting one against a limit written for the other is not, so state the method with the number. MEASURED

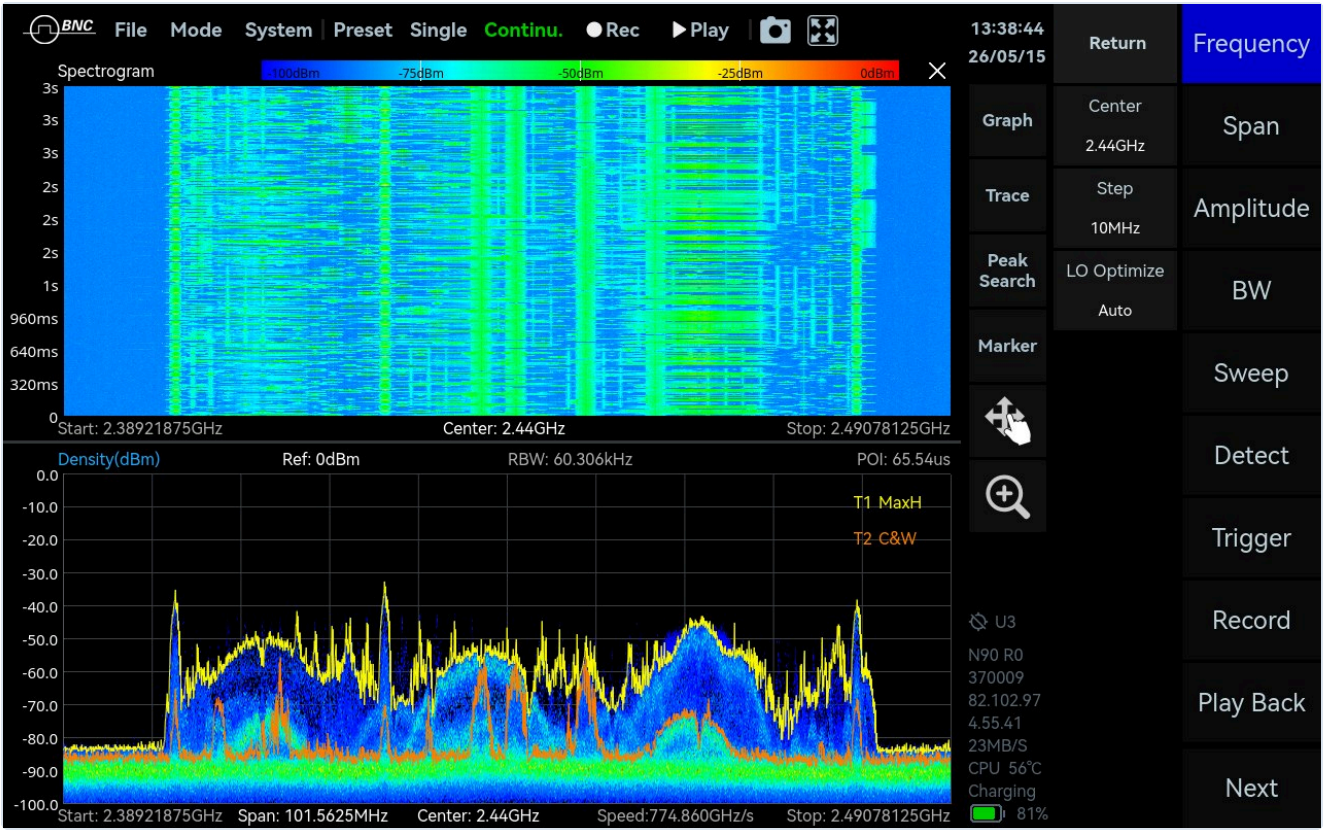

Figure 15. Real-time density with its spectrogram: 2.389 to 2.491 GHz in a 101.56 MHz span at 60.3 kHz resolution, with the display reporting a 65.54 microsecond probability of intercept. That figure is 2 × 4096 × 8 ns to the digits the display carries, so equation (8) is visible here as a reading at a transform size no datasheet quotes. **Density accumulates every transform rather than the peak of a sweep**, so an intermittent emitter resolves into a persistent trace with its duty cycle visible instead of an occasional spike. The span field is checkable too: 101.5625 MHz divided by the 30.5176 kHz bin spacing is exactly 3328 bins, which is 3328 of the transform's 4096. It therefore sits above the published 100 MHz analysis bandwidth and below the transform's full 125 MHz extent, which makes it a display span rather than a usable-bandwidth claim. **MEASURED**

Figure 16. One acquisition presented four ways at once: power against time, max-hold and clear-write spectra, a two-second spectrogram, and I and Q across a 208 microsecond window with a sample-rate field reading 125 MHz. The span reads 2.3775 to 2.5025 GHz, which is 125 MHz: **the sample rate rather than the usable 100 MHz**, which is the confusion section 16 warns about, visible here as a setting. Read the spectrogram against the buffer arithmetic rather than past it: two seconds is far longer than the buffer holds at this rate, so that pane is a live rolling display and the IQ pane below it is the buffered case. **MEASURED**

Figure 13, Figure 14, Figure 15 and Figure 16 have one property in common that is worth naming as a purchasing test. Each reports a derived quantity **and** the raw quantities it was derived from, so a reader can check the derivation without trusting the device. Apply that test to any instrument you are evaluating: ask it for a measurement, then ask whether the screen contains enough to reproduce it. A device that reports only conclusions is asking for a kind of trust that a defensible measurement cannot be built on.

44 One engine, three form factors

A point from Part II applies to this family specifically, and it can be shown rather than asserted: within a product line, the form factor should change what the instrument survives and where it goes, not what it measures.

Figure 17. A product render of the rugged tablet, not a photograph, and the screen it carries is the capture already shown as Figure 8 rather than a second measurement: the timestamp, the spans, the temperature and the battery state all match it exactly. It is here to show the form factor, and the bezel reading ICX-400R, 9 kHz to 40 GHz, is the only thing in the image offered as evidence of anything. **The published interface list is identical across all three form factors**, the same languages, the same SCPI and the same measurement set, which is what makes a routine written against a module on a bench portable to this one. That the datasheets agree is published; that any given routine moves untouched follows from it as a reasonable inference, and is worth confirming for your own routine before you plan around it. It is the practical content of buying a line rather than a product. SCHEMATIC

Figure 17 makes the commercial case the way this handbook prefers to make it, which is by naming a property a buyer can test. Ask any vendor offering a family whether a routine written against one member runs unmodified against the others, and what specifically differs. **A family whose members need different code is three products sharing a brand**, and the difference shows up as integration cost long after the purchase.

The published physical differences are real and are the point of having three. Module weights start below 305 g for the 9.5 GHz USB variant and reach 665 g for the LAN version, against 1.5 kg for the handheld with its display and battery. Environmental classes are options rather than variants: a base class covering 0 to 50 °C, option 40 widening it to −20 to +65 °C and option 41 to −40 to +65 °C, published on the rugged and module datasheets.

45 Where this is not the answer

A handbook that only lists capability is a brochure. These are the cases where an ICX-FieldHawk is the wrong purchase, and naming them is how the rest of the document earns its credibility.

- **You need transmit.** This is a receiver family. Vector signal generation is not an ICX capability; the only published generation feature is a built-in signal generator covering 100 kHz to 6.3 GHz on USB models as option 02. A record-and-replay bench pairs an ICX capture with a separate generator, and BNC-AN-105 keeps the two halves apart throughout.
- **You need continuous wideband recording.** Sustained capture is published at 25 MHz, not 100 MHz. If the application genuinely needs minutes of unbroken 100 MHz recording, the right answer is a recorder with a storage array behind it.
- **You need DIFI or VITA 49 transport.** Not published. See section 23.

- **You are learning, teaching or prototyping on a budget, below 6 GHz.** The development class exists for exactly this, it is excellent at it, and an instrument bought for a job that never needed traceability is money spent on a column nobody will read.
- **You need a coherent array today.** Multiple-unit operation is a published function, but if the requirement is phase-coherent direction finding, ask what is specified about the relationship between units before committing, exactly as section 18 advises asking any vendor.

46 An uncertainty budget, worked

A bound like plus or minus 2.0 dB is the instrument's contribution. It is not the uncertainty of your measurement, and the difference between those two things is where most disputes about a number begin. Building the budget is short arithmetic and it is worth doing once so the shape is familiar.

Each contribution enters as a variance. A bound quoted as plus or minus *a* with no distribution stated is treated as rectangular, which contributes:

$$u^2 = a^2 / 3 \qquad (10)$$

where

a is the half-width of the stated bound, in decibels

u is the standard uncertainty contributed, in decibels

Sum the variances, take the root, and multiply by a coverage factor of two for approximately 95 percent confidence. Table 15 does that for a field measurement of the kind BNC-AN-107 works through in full.

Table 15. A worked budget, and it has two lessons that pull in opposite directions. **The instrument dominates:** 1.333 of a combined 2.032 is 66 percent of the variance, so the receiver bound really is the term to attack first, and halving it from ±2.0 to ±1.0 dB takes the expanded figure from 2.85 dB to 2.03 dB. **And it is not the whole story:** with a hypothetically perfect receiver the remaining terms still leave 1.67 dB expanded, so the antenna factor, the cable and the placement set a floor no instrument purchase gets you under. Budget both ends before deciding where the money goes.

Contribution	Stated bound	Variance contribution, dB^2
Instrument amplitude accuracy	±2.0 dB	1.333
Antenna factor calibration	±1.0 dB	0.333
Cable assembly loss	±0.6 dB	0.120
Mismatch	±0.5 dB, U-shaped	0.125
Repeatability of placement	±0.6 dB	0.120
Combined		**2.032**
Standard uncertainty		**1.425 dB**
Expanded, k = 2		**2.85 dB**

WHY ONE ROW USES A DIFFERENT DIVISOR

Every bound above is converted to a variance by dividing by three, which is what a rectangular distribution gives. **Mismatch is the exception and divides by two**, because mismatch uncertainty is U-shaped: the error arises from two reflections adding with a phase that is uniform, so the resulting amplitude spends most of its time near the extremes rather than near the middle. It is the standard treatment in radio-frequency metrology and it makes this row half again larger than a rectangular assumption would. Using the wrong divisor here is one of the commonest errors in a published budget.

WHAT A COVERAGE FACTOR OF TWO DOES AND DOES NOT PROMISE HERE

The convention that $k = 2$ corresponds to about 95 percent assumes the combined distribution is roughly normal, which it is when many comparable terms add. **In this budget it is not:** one rectangular term carries two thirds of the variance, and a rectangular distribution runs out of range at 1.73 standard deviations. So at $k = 2$ the true coverage here is above 95 percent rather than below it. The interval is conservative, which is the right direction for a bound you intend to defend, but **"approximately 95 percent" is a convention in this table rather than a computed figure**, and it is worth saying so before someone else does.

THE TWO THINGS THIS BUDGET TEACHES A SPECIFIER

First, **a receiver bound is a floor on your uncertainty, not a description of it**, so a requirement written only against the instrument has specified the smallest term. Second, the arithmetic only works if every contribution is stated as a bound with a distribution. A device that publishes no amplitude accuracy at all cannot enter this table, which is the practical sense in which such a device cannot be used for a defensible measurement: not that it is inaccurate, but that its accuracy is unstated and therefore unbudgetable.

PART V

Choosing, and reference

A selection method that works against any vendor, and the numbers with their conditions.

47 Bandwidth, and the four numbers a span can mean

Section 8 separated sample rate from bandwidth. In practice a span quoted on a screen or a datasheet can mean any of four things, and confusing them accounts for a surprising share of the disagreements between a measurement and a specification.

Table 16 separates them.

Table 16. Four meanings of a span, in decreasing order of size and increasing order of usefulness. **The fourth is the one that governs a measurement and the first is the one that sells the product.** A vendor who publishes the second and fourth has told you something; one who publishes only the first and second has left the question open.

The number	What it is	How to recognize it
Tuning range	Where the tuner can be pointed	The widest number on the datasheet, and never a performance claim
Sample rate	Samples per second crossing the interface	Often displayed as a span, because a complex rate represents a span of the same width
Instantaneous bandwidth	What is digitized at once	Should be smaller than the sample rate. If they are equal, the datasheet is quoting the rate
Usable bandwidth	What is flat and trustworthy	Smaller again, by the filter transition. Around 0.8 of the rate is typical, and it is the only one of the four you can actually measure across

A WORKED EXAMPLE OF THE CONFUSION

A screen reporting a 125 MHz span at 125 MSPS on an instrument whose published analysis bandwidth is 100 MHz is not contradicting itself. The span field is reporting the sample rate, which represents 125 MHz of complex spectrum; the published 100 MHz is what the anti-alias and decimation filters leave flat. **Both numbers are correct and they answer different questions**, and an operator who records the first as the measurement bandwidth has overstated their coverage by 25 percent without doing anything wrong on the glass.

48 Nyquist zones, and deliberate undersampling

Sampling at rate f_s divides the frequency axis into zones $f_s/2$ wide. The first runs from zero to $f_s/2$, the second from $f_s/2$ to f_s, and so on. **Everything in every zone lands on the same output band**. That is the fact aliasing is usually taught as a hazard.

Used deliberately it is a technique rather than a hazard. If the analog front end passes only one zone, sampling can be slower than the signal's frequency and still recover it completely, because the ambiguity the theorem warns about has been removed by the filter instead of by the sample rate. This is why a converter quoting an analog input bandwidth far above its sample rate is not contradicting itself. Two details matter in practice: **even-numbered zones come out frequency-reversed**, so the spectrum is a mirror image and any downstream processing has to know it, and the higher the zone the more the converter's aperture jitter costs you, so sensitivity falls off with zone number even though the arithmetic does not.

HOW TO READ A DIRECT-SAMPLING SPECIFICATION

A part quoting 32 gigasamples per second has its first zone from zero to 16 GHz. An input range quoted to 18 GHz therefore includes second-zone operation, and that is a design choice rather than a contradiction. **Ask which zone a figure applies to**, and ask for the sensitivity in that zone rather than the one in the first, because they are not the same number and only one of them is usually published.

49 Aliasing, and the spurs that are not there

One more artifact class deserves naming, because it produces the most confident wrong conclusions. A sampled system cannot distinguish a signal at frequency *f* from one at *f* plus any multiple of the sample rate. Everything outside the digitized span folds into it, and once folded it is indistinguishable from a signal that was really there.

That is what the anti-alias filter prevents, and it is why the bandwidth control of section 16 is not a convenience. Set a sample rate of 20 MSPS and leave the filter at 56 MHz and everything in that 56 MHz appears somewhere in your 20 MHz of data, at frequencies bearing no relation to where it actually was.

THREE ARTIFACTS, THREE TESTS, ONE HABIT

Alias: change the sample rate and see whether the feature moves. A real signal stays at its frequency; an alias moves, because its apparent position depends on the folding. **Image:** retune and see whether the feature stays symmetric about the new center.
Intermodulation product: add 10 dB of input attenuation. A calibrated display is re-referenced for its own attenuator, so a real signal does not move while an internally generated third-order product drops 20 dB. Pad ahead of the instrument instead, leaving the analyzer alone, and the real signal drops 10 dB while the product drops 30. Either way the discriminator is the same 20 dB of extra movement, because the product goes as the cube of the level at the mixer. Each of these is ten seconds of work and each settles a question that can otherwise consume a day.

The third test is worth a sentence more, because it is the one that catches a receiver being driven too hard. Intermodulation products are generated inside the receiver by strong signals, and they scale faster than the signals that produce them: a third-order product changes three decibels for every decibel of input. If padding the input by 10 dB drops a feature by 30 dB, the feature was never in the air. This is also the practical reason the input intercept figures of section 10 matter: they are the specification that tells you how much signal a receiver tolerates before it starts inventing.

50 Selection, by the job rather than by the datasheet

The method below is deliberately vendor-neutral. It produces a requirement that several suppliers can answer honestly, which is the only kind worth writing.

1. **Start with the shortest event you must not miss.** That fixes whether you need real-time capture at all, and if you do, equation (8) fixes the largest transform you may use.
2. **Then the widest span you must hold at once.** Instantaneous bandwidth, not tuning range. This is the number most often confused with the next one.
3. **Then the highest frequency.** Above roughly 6 GHz the integrated-transceiver class thins out and a different architecture is required, so this question can eliminate a whole tier of the market in one line.
4. **Then ask whether the answer has to be defensible.** If somebody outside your team must accept the number, you need a stated amplitude bound and a calibration date, and you have just moved into the fourth class of device.
5. **Then transmit, duplex and coherence**, each stated as a requirement about the relationship between channels rather than a count of them.
6. **Then timing.** Reference accuracy, holdover and timestamp accuracy, asked separately.
7. **Only then the software.** Which framework, which abstraction layer, and whether the processing investment survives a change of radio.

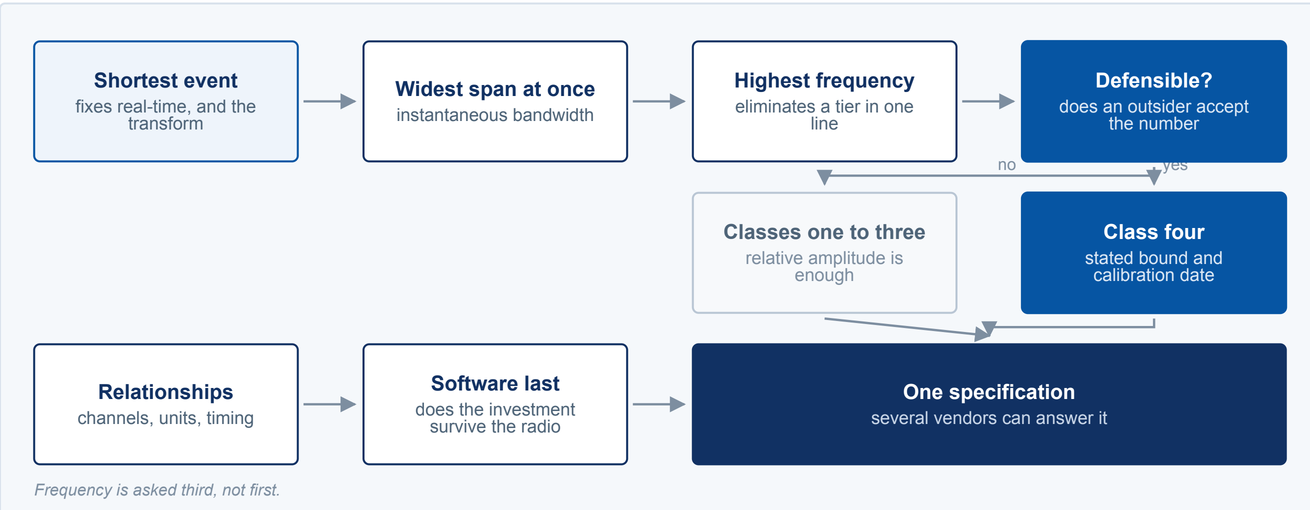

Figure 18. The order matters more than any single answer. Most requirements open with frequency; here it is the third question. **The fourth is the one that actually partitions the market**, because defensibility decides which class of device you are buying before any number is compared. Everything below the split is the same for both branches. SCHEMATIC

Figure 18 is the method as a picture; Table 17 turns it into a form you can put on a requisition.

Table 17. A vendor-neutral requirement set. Every row is written so that a supplier can answer it with a number or admit it does not publish one, which is the point: **the answers you cannot get are as informative as the ones you can.**

Ask for	Stated as	Why this wording
Shortest guaranteed event	microseconds, at 100 percent probability of intercept	Rules out any device that cannot state one
Instantaneous bandwidth	MHz digitized at once, and usable bandwidth separately	Separates the sample rate from the span
Sustained record rate	MHz sustained to host, distinct from burst	A single figure hides which one it is
Amplitude accuracy	±dB over a stated frequency range and conditions	The field whose absence defines the category
Reference	ppm disciplined, plus holdover over a stated interval	Holdover is the term distributed systems actually depend on
Timestamp	ns against 1PPS	Converts directly to meters at 0.2998 m/ns
Channel relationship	specified relative phase or delay between channels	A channel count is not a coherence claim
Interface	documented API, named languages, and transport format	Decides whether your software survives the hardware

51 Ten jobs, and what decides each one

Selection is easier when the job is named. Each entry below is a problem this class of instrument is bought for, the property that actually decides it, and the Berkeley Nucleonics application note where the arithmetic is derived rather than asserted.

Table 18 is that map.

Table 18. Ten jobs and what turns each one. Read the middle column as the question to ask a vendor, because it is the property that decides the outcome and it is rarely the property on the front of a datasheet. **Not one of these rows is decided by frequency range**, which is the specification buyers compare first.

Job	What actually decides it	Derived in
Programming an analyzer as a platform	Whether calibration crosses the software boundary with the samples	BNC-AN-101
Wide-area monitoring with automated classification	The capture, not the classifier. No model recovers what the front end discarded	BNC-AN-102
Detecting unmanned aircraft and their operators	Holding the control band whole. A hopping link occupies a channel for milliseconds	BNC-AN-103
Commissioning a satellite earth station	The archived baseline, the only timestamped evidence of a compliant state	BNC-AN-104
RF record and playback	The capture sets the ceiling; the generator only decides what is lost from it	BNC-AN-105
UAV-borne antenna measurement	Far-field geometry, and statistics over a population of pulses	BNC-AN-106
EMF safety and RF exposure	Which quantity the limit is written in, and the uncertainty budget behind it	BNC-AN-107
5G, Wi-Fi and Bluetooth coexistence	Holding a whole carrier edge to edge under one amplitude reference	BNC-AN-108
EMC pre-compliance on your own bench	Dwell time against the intercept relation, and the antenna factor arithmetic	BNC-AN-109
Designing a receiver into your own product	Time transfer and site geometry, not receiver noise figure	BNC-AN-110

The ten, in a sentence each

The table is the index. These are the arguments, one paragraph each, for a reader deciding which note to open.

- **Programming an analyzer as a platform.** The decision gets argued on noise figure and is actually about whether the samples are reachable and whether the calibration reaches them too. An uncorrected board hands a flowgraph a fraction of full scale; a calibrated instrument hands it a level in dBm with a stated bound.
- **Wide-area monitoring with automated classification.** The receiver has become a sensor feeding a statistical estimator, and the estimator inherits every property of the data. A swept receiver at 1 MHz resolution across 40 GHz with a 100 ms revisit dwells 2.5 microseconds per cell, which is one part in 40,000 of the revisit, so that is the fraction of any emitter's transmissions arriving while it is looking, whatever their rate. No model recovers what that discarded.
- **Detecting unmanned aircraft and their operators.** The control link is present in any one channel for milliseconds, and the operator is found by direction-finding the uplink rather than the aircraft. What a visiting receiver reports as empty is a function of when it looked.
- **Commissioning a satellite earth station.** The deliverable is the file the visit produces. An interference claim arrives months later as an assertion, and the commissioning archive is the only timestamped, instrument-traceable evidence of a compliant state.
- **RF record and playback.** A replay can never be cleaner than its capture: the analyzer's noise, spurs and imbalance are burned into the file and are indistinguishable from signal to the generator. The capture instrument sets the ceiling.
- **UAV-borne antenna measurement.** Far field begins near two aperture-squared over wavelength, and at exactly that distance the phase error across the aperture is one sixteenth of a wavelength. A pattern point is a statistic over a population of pulses, not a single reading.

- **EMF safety and RF exposure.** Limits are written in field strength or power density and an analyzer reports power at a connector, so the job is a conversion with an uncertainty budget behind it, of exactly the kind section 46 works through.
- **5G, Wi-Fi and Bluetooth coexistence.** Coexistence problems are between-standard problems, so an instrument that measures one standard at a time cannot see them. Holding 100 MHz gap-free puts the interferer and the victim in one acquisition against one amplitude reference.
- **EMC pre-compliance on your own bench.** No ICX specification table publishes an EMC measurement function, so this is ordinary spectrum work plus arithmetic you do yourself, not a compliance mode. The arithmetic that decides whether a scan is honest is dwell time against the intercept relation. Running a confirmation scan at search speed is the commonest way a pre-compliance result turns out optimistic.
- **Designing a receiver into your own product.** In a distributed radio-location node the receiver's own timing contributes a tiny fraction of the position-error variance while time transfer, inter-channel matching, enclosure temperature and site geometry contribute nearly all of it. One nanosecond is 0.2998 m, and geometry multiplies whatever survives.

52 Troubleshooting, by symptom

Table 19 is arranged by symptom, because a symptom is what you have when you open this page. The third column is the point of it: most of these symptoms have several plausible causes, and the discriminating check is what stops an afternoon going on the wrong one.

Table 19. Symptom to cause to a check that tells two causes apart. **The first two rows are the zero-IF signature from Part I**, and between them they account for a large share of the spurious-signal reports this class of hardware generates. The fifth row is the one to expect in software, because the spectrum stays healthy throughout and nothing announces itself.

Symptom	Probable cause	The check that discriminates	Fix
A carrier at the exact center of the span that moves with the tuning	DC offset or oscillator leakage in a zero-IF receiver	Retune by a known amount. A real signal moves in the display; this does not	Tune deliberately off the signal, or enable the receiver's offset-tuning mode
A mirror image at the negative of a signal's offset	Quadrature imbalance	Retune. The image stays symmetric about the new center	Run the device's IQ correction; accept the residual as an image-rejection figure
Spectrum looks empty	Tuned outside the signal, or a bandwidth filter far narrower than the sample rate	Widen the span and look for the carrier; read back the bandwidth setting	Retune; set bandwidth explicitly rather than trusting a default
Amplitude plausible but wrong	Automatic gain control on, or a reference level changed mid-campaign	Inject a known level from a calibrated source and compare	Turn device AGC off; fix and record the reference level
Intermittent decode failure, spectrum healthy	Overflow: samples discarded at the driver, leaving no gap in the delivered stream	Check the stream return code, not the console	Decimate at the source; cut display update rates; move per-sample work out of interpreted code
Weak signals disappear when a strong one is present	The strong signal is setting the gain and consuming the dynamic range	Remove the strong signal or filter it and see whether the weak one returns	Preselect or notch ahead of the receiver; reduce gain and accept the noise floor
Sensitivity worse than the datasheet	Cable loss ahead of the first amplifier, which is noise figure decibel for decibel	Measure the loss of the run, or move the receiver to the antenna	Amplify at the antenna, shorten the run, or budget for the loss knowingly
Two receivers disagree about phase	No common reference, or no specified relationship between channels	Lock both to one reference and repeat	Use a common reference and a hardware trigger; confirm what the vendor specifies about relative phase

53 Buy, build, or integrate

One decision sits above the selection method, and getting it wrong is more expensive than any subsequent choice. Three routes exist to a receiving system, and the arithmetic that separates them is not about radio performance.

Table 20 sets the three routes side by side.

Table 20. The three routes. **A build-versus-buy decision is settled by what the third row makes you own**, not by radio performance, and the radio design is the smallest item in it. The escape route most teams try first, buying a low-cost wideband board and correcting it in software, works on the bench and fails at the second unit, because the correction has to be established and maintained per unit rather than per design.

Route	What you own	What decides against it
Buy an instrument	The application and nothing below it	Unit cost at volume, and any requirement the instrument does not meet
Integrate a module	The enclosure, the thermal design, the conformity of the finished product, and the software above the interface	The mechanical and regulatory work, which is usually underestimated by more than the RF work
Build a receiver	The radio design, the calibration fixture, the traceability chain, qualification, and ten years of obsolescence engineering	**Volume.** None of the recurring costs scale down, so below a certain annual quantity the build route is more expensive at every point in its life

The crossover is a volume, and computing it beats arguing about it. A fixed development and qualification cost spread over a lifetime volume meets a flat unit price at one quantity, and BNC-AN-110 derives that crossover in full for an embedded monitoring product. The result most teams find surprising is how high the crossover sits once calibration and obsolescence engineering are counted honestly.

54 Compliance, export and supply chain

The regulatory work in section 53 is the item most often left out of a schedule, so it is worth naming what it consists of. Five separate regimes apply to a receiving product, they are administered by different bodies, and satisfying one says nothing about the others.

Table 21 separates them.

Table 21. The five regimes. **Only the first two are alternatives to each other**, since RED absorbs the EMC essential requirements for anything in its scope; the rest apply on top, in whatever combination fits the product and the market. The common failure is to treat a supplier's declaration for one as covering the others, and the second most common is to inherit a module's evaluation-kit exclusion into a product that is no longer an evaluation kit.

Regime	What it governs	What it means for a receiver
EU Radio Equipment Directive, 2014/53/EU	Equipment that intentionally emits **and/or receives** radio waves for radio communication or radiodetermination	Receive-only equipment is in scope on the face of the definition. **Custom-built evaluation kits are excluded**, which is the whole reason a development board can ship with paperwork a finished product cannot use
EU EMC Directive 2014/30/EU and the Low Voltage Directive	Emission and immunity; electrical safety	Applies where RED does not. RED itself incorporates the EMC essential requirements by reference, so a product in RED scope does not additionally carry the EMC Directive
FCC equipment authorization, 47 CFR Part 15	Radio-frequency devices sold in the United States	A receiver is a digital device and an unintentional radiator; anything that transmits needs authorization on its own account
EU Cyber Resilience Act, Regulation (EU) 2024/2847	Cybersecurity of products with digital elements, across their supported life	**A receiver with firmware, host software and a driver is in scope.** Manufacturer reporting of actively exploited vulnerabilities and severe incidents applies from 11 September 2026, on a 24-hour first notification; the main obligations follow on 11 December 2027
Export control (EAR, and ITAR where applicable)	Whether an item may leave the country, and to whom	Wideband receiving equipment attracts attention. **The classification is the manufacturer's to make and defend**, and a buyer should ask for it in writing rather than infer it

The evaluation-kit exclusion deserves the emphasis. A development board is exempt because it is a development board. Putting it in an enclosure and selling it makes it a product, and the product is assessed on its own, from scratch, by whoever puts their name on it. That work does not transfer from the board vendor and it is not small: it is a test campaign, a technical file, a declaration of conformity, and a person willing to sign it.

> **WHY THIS HANDBOOK DOES NOT PUBLISH A CLASSIFICATION**
>
> **A specific export classification is a legal determination, and a marketing document is the wrong place to publish one.** Ask for it in writing, against your destination and end use, from whoever is shipping the item. Any supplier that answers the question casually in a brochure is telling you something about how they handle the rest of their paperwork.

The Cyber Resilience Act is the newest of the five and the one most likely to be missed, because it does not look like radio regulation. It reaches software, and a receiver is a software product with an antenna on it. Two features make it consequential for a buyer. The reporting duty applies to products **already on the market**, so it is not something a supplier can defer until the next design. And it runs for the whole supported life of the product, which turns a question most datasheets never answer, how long this thing will be supported, into a regulated commitment. Ask for the support window in writing, and ask who issues firmware updates and how they are signed.

55 Provenance, and the questions procurement will ask

Where a receiver was designed and built has become a purchasing criterion in its own right, independent of performance. Under **Section 889 of the FY2019 National Defense Authorization Act**, federal agencies may not procure covered telecommunications equipment as a substantial or essential component of any system. Part A, the procurement prohibition, took effect on August 13, 2019; Part B, which bars agencies from contracting with an entity that *uses* such equipment, took effect on August 13, 2020. Part B is the one that surprises people, because it reaches past what you are selling to what you are running.

The practical consequence for a receiver is that the answer has to hold at component level, not at brand level. A platform assembled from modules is only as clean as its least documented module, and a bill of materials nobody has read is not an answer. Three questions establish more than any assurance does:

- **Who designed the radio, and where is it manufactured?** Not where the box is assembled, which is a different question and often a different country.
- **Can you produce a bill of materials to component level, and will you flag changes?** A supplier who cannot has not been asked before.
- **What is the support and repair path, and who performs the calibration?** An instrument returned across a border for recalibration is a supply-chain fact, not a service detail.

So that we are held to the same standard: **Berkeley Nucleonics is a California company**, and the ICX line is supported and calibrated through it. That answers the third question and part of the first. **It is deliberately not an answer to component-level origin, which is a bill-of-materials question and belongs in writing against a specific configuration rather than in a handbook.** Ask us for it the same way you would ask anyone, and treat a brochure paragraph from any supplier, including this one, as the start of the conversation rather than the end of it. Nothing here is a claim about anybody else's provenance.

56 Reading somebody else's datasheet honestly

This handbook cannot compare products for you, and would not be worth much if it did, because the comparison you need is against your requirement rather than against a league table. What it can do is give you the questions that make any datasheet answer for itself.

- **Find what condition each number carries, and treat an unconditioned number as absent.** A noise figure without a frequency, a sweep speed without a resolution bandwidth, a phase noise without an offset: each is a number that cannot be reproduced and therefore cannot be checked. This applies to our datasheets exactly as much as to anyone's.
- **Look for the pairs that should both be present.** Sample rate and usable bandwidth. Burst and sustained. Resolution and effective bits. Preamp-on and preamp-off damage limits. Where only the flattering half of a pair appears, that is a choice somebody made.
- **Check whether a range is a specification or a coverage claim.** Many devices publish a frequency range far wider than the range over which any performance figure is stated. The question to ask is not how wide the range is but over what part of it the accuracy holds.
- **Count the tiers with specification tables.** A family advertised in five frequency tiers with published tables for two has told you something. This is worth applying to the ICX-FieldHawk too: **only the 9.5 GHz and 40 GHz tiers carry published tables**, and nothing in this handbook is written against any other.
- **Ask what is specified about relationships, not just about parts.** Between channels, between units, between the calibration and the samples. Section 18 lists these and they are consistently the least published and the most consequential.

THE SINGLE MOST USEFUL QUESTION TO SEND A VENDOR

Is device calibration applied before the samples cross your interface, and what bound do you hold to on the samples I receive? It is one sentence, it has a yes-or-no first half and a numeric second half, and it cannot be answered evasively without the evasion being obvious. Almost everything else in a selection follows from the answer.

57 Measurement practice that applies whatever you bought

1. **Let it warm up.** Published accuracy figures assume it. Measurements from a cold start are covered by no specification on any datasheet.
2. **Set the reference level deliberately.** Amplitude accuracy, intermodulation performance and displayed noise floor all move with it, and an automatic reference level during a comparison invalidates the comparison.
3. **Read back every setting you request.** A sample rate, a bandwidth or a frequency is a request, not an assignment: hardware has step sizes, and a filter designed against the value you asked for and fed the value you got produces a slow error nothing downstream catches.
4. **Record the state, not just the trace.** Center frequency, span, sample rate as read back, resolution bandwidth and video bandwidth, decimation, transform size, window, detector, trace mode, reference level, attenuation, preamplifier state, antenna and its factor, cable assembly, and a timestamp. A capture without that sidecar is a pile of integers.
5. **Fix the cable and leave it fixed.** Loss in flexible assemblies changes when the cable moves, so swapping one mid-campaign is an uncalibrated step in the middle of the record.
6. **Record the resolution bandwidth with every noise figure.** A displayed average noise level per hertz becomes a real floor only after ten times the log of the bandwidth is added, and the two get confused constantly.

58 Two worked selections

The method is easier to trust when it has been run. These are two real shapes of requirement, worked through the questions of section 50 to a class of device rather than to a product.

A university teaching and research laboratory

The job is teaching modulation and protocol work to students, plus occasional research captures below 6 GHz that end up in papers. Budget is per-seat rather than per-instrument.

- **Shortest guaranteed event:** not a constraint. The signals are generated in the lab and repeat on demand, so intercept probability does not decide anything.
- **Widest span held at once:** 20 MHz covers every waveform on the syllabus.
- **Highest frequency:** 6 GHz. Everything taught lives below it.
- **Does the answer have to be defensible?** For teaching, no. **For the research captures that reach a paper, yes**, and this is the split that matters.
- **Conclusion:** development-class transceivers for the teaching seats, where the open ecosystem and the price are exactly right, and **one** instrument-grade receiver shared across the group for anything that gets published. Buying instruments for every seat would waste money on a column nobody grades; buying none would mean every published amplitude rests on an uncalibrated board.

A distributed monitoring network across a site

The job is unattended nodes across a large industrial site, reporting occupancy and flagging unexpected emitters, expected to run for years and to be extended later.

- **Shortest guaranteed event:** short. The emitters of interest are intermittent, which makes intercept probability the first specification rather than an afterthought, and equation (8) then fixes the largest transform each node may use.
- **Widest span held at once:** as wide as the site's allocation, held continuously rather than swept, for the reason Figure 8 shows.
- **Highest frequency:** set by the site, and if it reaches past 6 GHz that decision has already eliminated a tier of the market.
- **Defensible?** Yes. An occupancy claim that leads to an enforcement conversation has to survive being

questioned.

- **Timing:** the deciding requirement, and the one most likely to be under-specified. Ask for holdover separately, because the nodes will lose their discipline.
- **Transport:** name a format. This system is explicitly expected to outlive its hardware, which is the case section 23 was written for.
- **Conclusion:** instrument-grade modules with a published intercept relation, a stated amplitude bound, a timing option and a documented interface. **And a written answer on transport format before committing**, since that is the requirement most likely to be unmet, including by us.

> WHAT BOTH EXAMPLES HAVE IN COMMON
>
> Neither was decided by frequency range, which is the specification most buyers compare first, and **both split at the question of whether the answer has to be defensible**. That question sorts this market more cleanly than any performance figure, and it is not on any datasheet because it is a property of your project rather than of the hardware.

59 What lives where: a working band reference

Half of the work in an unfamiliar capture is knowing what you are probably looking at before you identify it. The allocations below are the ones that account for most of what a wideband receiver encounters below 8 GHz, and carrying them in your head beats looking them up each time.

Table 22 is the short version.

Table 22. What a receiver meets below about 8 GHz. The last row is the argument of section 5 in one line: a whole class of platform cannot see a band that is now in ordinary consumer use.

Range	What is there	Why a receiver cares
300 to 512 MHz	Land mobile, public safety, telemetry	Narrowband, dense, and channelized; resolution matters more than span
824 to 960 MHz	Cellular 850/900, and the US 902 to 928 MHz ISM band that carries LoRa, telemetry and much of the low-cost IoT world	The 902 to 928 MHz allocation is Region 2; Europe's equivalent is 863 to 870 MHz, so a device tested in one place is not necessarily legal in the other
978 and 1090 MHz	UAT and ADS-B	Short pulsed bursts, always on, and a reliable signal of a working receive chain
1164 to 1610 MHz	**Three different things, not one.** The navigation sub-bands proper: GPS L5 at 1176.45 MHz, L2 at 1227.60 MHz, L1 at 1575.42 MHz, plus GLONASS, Galileo and BeiDou. Between them, 1215 to 1400 MHz carries long-range air surveillance radar at megawatt peak power, and 1400 to 1427 MHz is a passive band in which all emission is prohibited	**In the navigation sub-bands** the signals are far below the noise floor by design, so anything visible without processing gain is interference rather than a satellite. That rule does not extend across the middle of the range, where the strongest signal a wideband receiver meets all day is entirely legitimate
1700 to 2200 MHz	AWS, PCS and the cellular uplink and downlink pairs, and since 2024 also satellite direct-to-device operating in that same terrestrial spectrum	Duplex spacing makes the pairing recognizable before anything is decoded. **An emitter here is no longer necessarily on the ground**, which section 61 takes up
2400 to 2500 MHz	The 2.4 GHz ISM band: Wi-Fi, Bluetooth, and most of the consumer radio world	The most congested band a field receiver meets, and the reason max hold exists
3550 to 3700 MHz	CBRS, shared under a three-tier authorization framework	Incumbent radar, licensed priority access and general authorized access share one 150 MHz band, so occupancy here changes with location and time by design
3700 to 4140 MHz	C-band. The lower part was cleared and auctioned for 5G; in July 2026 the FCC adopted an order releasing a further 160 MHz at 3.98 to 4.14 GHz, making a largely contiguous 440 MHz band	**The band is still moving as this is written.** It also runs up against radio altimeter spectrum, which is why out-of-band emission is scrutinized here more than almost anywhere else
5150 to 5895 MHz	The U-NII bands: Wi-Fi 5 GHz, and 5.8 GHz ISM	Radar detection is mandatory in the U-NII-2 sub-bands, so this range contains equipment actively listening as well as transmitting

Range	What is there	Why a receiver cares
5925 to 7125 MHz	6 GHz unlicensed in the United States	**The band where the common integrated-transceiver silicon runs out.** A platform built on it reaches 6.0 GHz and stops

Above 8 GHz the naming changes from allocations to the IEEE letter bands, and those are worth knowing because radar, satellite and point-to-point work is described in them almost exclusively.

Table 23 gives the letters.

Table 23. The IEEE letter bands, per IEEE Std 521. **Everything from X band upward is out of reach of the 6 GHz silicon class entirely**, which is not a criticism of it; it is a statement about which problems it was built for. Of the two ICX tiers, the 9.5 GHz variant reaches into X band and the 40 GHz variant covers the table.

Letter	Range	Typical occupants
L	1 to 2 GHz	Long-range surveillance radar, GNSS, some telemetry
S	2 to 4 GHz	Weather and airport radar, 2.4 GHz ISM, CBRS, C-band 5G
C	4 to 8 GHz	Satellite downlink, weather radar, 6 GHz unlicensed
X	8 to 12 GHz	Marine and airborne radar, satellite uplink
Ku	12 to 18 GHz	Satellite television and broadband downlink
K	18 to 27 GHz	Automotive radar at 24 GHz, satellite
Ka	27 to 40 GHz	High-throughput satellite, 5G millimeter wave, automotive at 77 GHz above this band

60 Where the bands are going

A reference that only describes today ages badly, and a receiver bought today will be in service for years. Two movements are worth writing into a specification now.

The upper mid-band. 3GPP calls 7.125 to 24.25 GHz FR3, and it is the most active piece of spectrum policy in the world at the moment. Resolution 256 (WRC-23) set studies running on 4400 to 4800 MHz, 7125 to 8400 MHz or parts of it, and 14.8 to 15.35 GHz, and those studies feed WRC-27 agenda item 1.7, which is where identification for mobile use would actually be decided. **WRC-23 scheduled the question; it did not answer it.** It is the leading candidate spectrum for 6G because it sits between the coverage of the sub-6 bands and the capacity of millimeter wave.

> READ THE PREVIOUS TWO TABLES TOGETHER WITH THAT PARAGRAPH
>
> **The whole of FR3 sits above the 6.0 GHz transceiver ceiling**, with 1.125 GHz of clear air between them, so the band the next decade of cellular work happens in is out of reach of the cluster most of this market is built on. It is not out of reach of the other cluster: the direct-sampling parts of section 6 already publish inputs into the teens of gigahertz, so **the upper mid-band is a large part of why that cluster exists**. For an instrument buyer the consequence is simpler than the silicon story. FR3 begins inside C band, which runs to 8 GHz in the letter scheme, and continues through X, Ku and into K. This is not an argument that one product is better than another; it is an argument that a frequency requirement written against a 2026 roadmap eliminates a whole architecture before any other specification is considered. Of the two ICX tiers, the 40 GHz variant covers FR3 entirely and the 9.5 GHz variant reaches only its lowest 2.4 GHz.

Sharing rather than allocation. The second movement is structural. CBRS put three tiers of user in one band under an automated coordinator, and 6 GHz put unlicensed devices under automated frequency coordination in the presence of incumbent links. Both replace a static allocation with a database that changes who may transmit, where, and when. For a receiver this changes the job: **occupancy becomes a function of time and place by**

design, not by accident, so a single visit to a band no longer characterizes it even in principle, and evidence about who was transmitting has to carry a timestamp and a position to mean anything at all.

61 Two more movements a receiver will meet

Signals in the cellular bands may now be coming from orbit. 3GPP Release 17 brought non-terrestrial networks into the same air interface as terrestrial cellular, and in 2024 the United States adopted a supplemental-coverage-from-space framework under which satellite operators transmit in a mobile operator's own terrestrial spectrum, as a secondary user, to reach ordinary handsets. Constellations are doing this today in PCS and AWS spectrum.

> WHAT THAT DOES TO THE BAND REFERENCE TWO PAGES AGO
>
> The row for 1700 to 2200 MHz used to mean terrestrial base stations and handsets. It now also means satellites, at very different power, geometry and Doppler. **For a monitoring job this breaks an assumption that used to be safe**, which is that an emitter in a cellular band is on the ground and can in principle be walked to. Direction finding against a source moving at orbital velocity is a different problem, and an interference case that assumes a terrestrial source can now be wrong about the sky.

And the fronthaul argument of section 23 has a second front. Open radio access network work disaggregates a base station into a radio unit and processing units from potentially different vendors, connected by an open fronthaul interface that carries digitized radio across an Ethernet network. It is the same structural story this handbook told about digital IF transport, in a much larger market: once an analog interface becomes packets, interoperability stops being automatic and has to be specified, tested and certified. The practical consequence for test equipment is that **the thing under test is increasingly a packet stream rather than a cable**, and an instrument that can only measure at an antenna port measures one half of a modern radio site.

62 Quick reference

Table 24 is the page to photocopy.

Table 24. Every governing relation and published value in this handbook, each with the condition that makes it true. Nothing here is quotable without the third column.

Quantity	Relation or value	Condition
Guaranteed intercept	POI = 2 N D × 8 ns	100 percent, full amplitude
Frame rate	10^9 / (N D × 8) frames/s	gap-free engine
Bin spacing	125 MSPS / (N D)	125 MSPS is the engine's undecimated rate; the delivered complex rate is 125 MSPS / D
POI × bin spacing	exactly 2	at every setting
Usable bandwidth	about 0.8 of sample rate	filter transition band, typical
Complex Nyquist	bandwidth B needs complex rate B	not 2B
Noise floor in RBW	DANL + 10 $\log_{10}$(RBW in Hz)	−99.9 dBm at −159.9 dBm/Hz in 1 MHz
Range from timing	0.2998 m per nanosecond	before geometry
Reference offset	f_c × ppm × 10^{-6}	1 ppm at 40 GHz is 40 kHz
Integrated transceiver span	70 MHz to 6.0 GHz receive; 47 MHz to 6.0 GHz transmit	AD9361 published range; the ceiling most of the market inherits
ICX amplitude accuracy	±2.0 dB / ±3.0 dB	to 9.5 GHz / 9.5 to 40 GHz

ICX values: the handheld, rugged and USB datasheets, 10-minute warm-up, 25 °C, spur reject standard. Transceiver span: Analog Devices published product data.

Quantity	Relation or value	Condition
ICX DANL	−159.9 / −167.5 dBm/Hz	1 GHz, RBW 1 kHz, 400 / 090 tier
ICX phase noise	−107.5 / −101.6 dBc/Hz	1 GHz, 10 kHz offset
ICX IQ capture	125 MSPS, 128 Mbyte, decimation 1 to 4096	100 MHz burst, 25 MHz sustained

ICX values: the handheld, rugged and USB datasheets, 10-minute warm-up, 25 °C, spur reject standard. Transceiver span: Analog Devices published product data.

63 A specification you can send to several vendors

The method in section 50 produces requirements. This section turns them into wording, because a requirement phrased badly gets answered by everyone and distinguishes nobody. Each line below is written so that a supplier must either give a number or say it does not publish one.

A vendor-neutral requirement form. **Line 5.4 is the one that sorts the market**, and it is a yes-or-no question that no datasheet field answers directly. Line 3.2 is the second most useful, because a supplier who can answer holdover has measured something most of the market does not.

```
1  ACQUISITION
1.1  Instantaneous bandwidth digitized at
one instant:        ____ MHz
1.2  Usable bandwidth within that span,
stated separately:    ____ MHz
1.3  Guaranteed intercept at 100%
probability, with the
     transform size it assumes:
____ us at N = ____
1.4  Sustained record rate to host,
distinct from burst:      ____ MHz
sustained

____ MHz burst

2  AMPLITUDE
2.1  Absolute amplitude accuracy, as a
bound, with the
     frequency range and conditions it
holds over:            +/- ____ dB
2.2  Displayed average noise level, with
frequency and RBW:   ____ dBm/Hz
2.3  Maximum input, preamp off and preamp
on:                 ____ / ____ dBm
2.4  Input intercept, third order, with
reference level:      ____ dBm

3  TIMING
3.1  Reference accuracy, disciplined:
____ ppm
3.2  Holdover over a stated interval after
discipline is
     removed:
____ ppm over ____
3.3  Timestamp accuracy against 1PPS:
____ ns

4  MULTI-CHANNEL  (omit if single channel)
4.1  Specified relative phase or delay
between channels:      ____
4.2  Common reference and hardware trigger
provided:          yes / no

5  INTERFACE
5.1  Documented programming interface,
languages named:       ____
5.2  Transport format, and whether VITA 49
or DIFI:          ____
```

```
5.3  Driver for an open hardware
abstraction layer:           yes / no
5.4  Is device calibration applied before
samples cross
     the interface:
yes / no
```

Two notes on using it. Send the same form to everyone, including us, and compare the shape of the blanks rather than the numbers: a supplier who leaves four lines empty has told you which class of device you are being offered. And **do not weight the lines before the answers come back**, because the temptation afterwards is to weight them toward whoever answered best.

64 Standards worth knowing by number

Table 25 collects the standards this handbook refers to, with what each one governs. A specification that cites a standard by number is answerable; one that describes a capability in prose is not.

Table 25. Standards by number, and what each is for. **Cite these in a requirement rather than describing the capability**, because a number is answerable and a description is not. Note the second row in particular: naming DIFI in a specification is the difference between a receiver you can replace and one you cannot.

Standard	What it governs	Where it bites
ANSI/VITA 49	Carrying digitized RF with its metadata over a network	Flexible enough that two compliant implementations may still not interoperate
IEEE-ISTO Std 4900-2021 (DIFI)	A constrained profile of VITA 49.2 that defines what is required	The one to name if a distributed architecture must outlive its hardware
IEEE Std 1139	Vocabulary for frequency and time metrology	Where accuracy and stability are separated precisely, in terms of the Allan deviation and its relatives
ITU-T G.810	Definitions and terminology for synchronization networks	**The source of holdover, free-run and locked as defined terms.** This is the number to cite when a supplier disputes what holdover means
IEEE Std 488.2 and SCPI	Instrument command language	The path that lets an existing test rack drive a new instrument without new software
CISPR 16-1-1	Radio disturbance measuring apparatus	Sets the receiver requirements behind any pre-compliance claim
JCGM 100:2008 (GUM)	Expression of uncertainty in measurement	The source of the rectangular-distribution treatment in section 46
ITU-R SM.1794 and the ITU-R Handbook on Spectrum Monitoring	International practice for spectrum monitoring	The framework a regulator will expect a monitoring result to sit inside

65 A reproducibility checklist

Whatever class of device you bought, a capture that cannot be reproduced is a capture that will be argued about. Copy this into a log and fill it in. **A capture without these entries is not reproducible, whatever else is true of it.**

1. Instrument model and serial number, and its calibration date if it has one.
2. Ambient temperature, and warm-up elapsed before the first reading.
3. Center frequency, span, resolution bandwidth and video bandwidth.
4. Sample rate **as read back from the device**, not as requested.
5. Decimation, transform size, window, detector and trace mode.
6. Reference level, attenuation and preamplifier state.
7. Whether device automatic gain control was on or off, stated explicitly.
8. Antenna, its antenna factor at the frequency of interest, and the cable assembly used.
9. Software versions: framework, abstraction layer, and every out-of-tree module.
10. Overflow or dropped-sample events with timestamps, or an explicit statement that none occurred.
11. The result of an amplitude check against a known level, with the date it was taken.

> **WHY THE FOURTH LINE IS THE ONE PEOPLE SKIP**
>
> A sample rate, a bandwidth and a frequency are **requests**, not assignments. Hardware has step sizes and clock divider trees, and a device asked for 20.000000 MSPS may deliver 19.999999 or 19.200000. Every downstream filter designed against the value you asked for and fed the value you got carries a slow frequency error that survives every other check in the chain. Read back, and record what came back.

66 Definitions, symbols and further reading

Aliasing
The folding of signals outside the digitized span into it. Once folded they are indistinguishable from signals that were really there, which is what the anti-alias filter exists to prevent.

Antenna factor
The conversion, in dB per meter, between the field at the antenna and the voltage at its connector. Required before any receiver reading becomes a field strength.

Back pressure
The propagation of a stall upstream through a processing graph when a block cannot consume as fast as its input arrives.

Coverage factor
The multiplier applied to a standard uncertainty to give a stated confidence. A factor of two is conventionally quoted as about 95 percent, exactly so only for a near-normal combined distribution.

Detector
The rule that reduces the many samples behind one displayed point to a single value: peak, sample, average or RMS. Only an RMS detector, or an average detector that averages power rather than decibels, reports correct power for a noise-like signal; log averaging under-reads noise by 2.51 dB.

Digital down-conversion
Shifting a band to zero frequency by multiplying the samples against a numerically generated oscillator, filtering, and only then decimating. The arithmetic equivalent of a mixer, and what a published decimation range specifies.

Dilution of precision
How much the geometry of a set of sensors or satellites multiplies a timing error into a position error. A value of three turns 10 m of range uncertainty into about 30 m of position uncertainty.

Direct sampling
Digitizing the antenna signal without analog down-conversion. Selectivity becomes arithmetic, and the converter sets the frequency ceiling.

Effective number of bits
Converter resolution derived from measured signal to noise and distortion, rather than from the nominal bit count. Always lower.

Equivalent noise bandwidth
The width of the ideal rectangular filter that would pass the same noise power as the window in use. It converts bin spacing into resolution bandwidth, and it is 1.9761 for the Blackman-Nuttall window and about 3.77 for a flat-top.

Error vector magnitude
The distance between where a received symbol landed and where it should have landed, as a fraction of the constellation. It measures modulation quality and is set by noise, phase noise and loop bandwidths rather than by calibration.

Gap-free
Every input sample reaches a transform, with no dead time for retune or processing. A property of the acquisition, not of the display.

Holdover
How far a reference drifts once external discipline is removed. Governed by aging and temperature coefficient, and usually unpublished.

Instantaneous bandwidth
The span digitized at one instant. Distinct from tuning range, which is covered by moving the tuner.

Intercept, guaranteed
The observation length over which an event is captured at full amplitude with certainty. A property of a real-time engine, and unrelated to the distortion sense of the word below.

Intercept, third-order
The extrapolated input level at which a receiver's third-order distortion products would equal the signals producing them. A measure of how much signal the receiver tolerates before it invents one.

Intermodulation product
A spurious signal generated inside a receiver by two or more strong inputs. A third-order product changes three decibels for every decibel of input, which is the test that identifies it.

Minimum detectable signal

The noise floor referred to the input: thermal noise density at 290 K, plus the noise figure, plus ten times the log of the measurement bandwidth.

Preselection
Filtering ahead of the receiver so that emitters outside the band of interest do not consume its dynamic range.

Probability of intercept
See Intercept, guaranteed. The two names are used for the same quantity, and the guarantee form is the one a specification can be written against.

Quadrature imbalance
A gain or phase mismatch between the in-phase and quadrature paths of a zero-IF receiver, seen as an image at the negative of a signal's offset.

Rectangular distribution
The assumption applied to a bound quoted without a distribution, contributing the square of its half-width over three to a variance budget.

Resolution bandwidth
The width of the filter or transform bin each displayed point represents. It sets the noise floor, so an amplitude quoted without it is not a measurement.

Spur reject
An instrument mode that suppresses internally generated spurious responses, at a cost in sweep speed. **It is a standing condition on every published ICX figure quoted here**, so the same instrument with it bypassed is quoting a different number.

Spurious-free dynamic range
The range between a signal and the largest spurious response the receiver generates from it. Distinct from noise-limited dynamic range.

TCXO and OCXO
Temperature-compensated and oven-controlled crystal oscillators. The oven costs power and warm-up time and buys roughly an order of magnitude in stability and considerably more in holdover, which is why it is an option rather than the standard fit.

Traceability
The unbroken chain from a reading to a national or international standard, with a stated uncertainty at each step. What separates a measurement from a reading.

U-shaped distribution
The distribution of a mismatch error, whose phase is uniform so its amplitude concentrates at the extremes. It contributes the square of its half-width over two, half again more than a rectangular assumption would.

Zero IF
Mixing the wanted band directly to baseband as I and Q. The architecture behind most integrated transceivers, and the source of three characteristic impairments.

Symbols used in this paper.

Symbol	Meaning	Units
N	transform size	points
D	decimation factor	dimensionless
f_s	complex sample rate	samples per second
f_c	carrier or tuned frequency	hertz
POI	probability of intercept	seconds

ACPR adjacent channel power ratio
ADC analog to digital converter
CRA Cyber Resilience Act
DANL displayed average noise level
DIFI Digital IF Interoperability
ENBW equivalent noise bandwidth
ENOB effective number of bits
IIP3 input third-order intercept point
IQ in-phase and quadrature
MDS minimum detectable signal
EVM error vector magnitude
NF noise figure
OCXO oven-controlled crystal oscillator
OFDM orthogonal frequency-division multiplexing
P1dB 1 dB compression point
POI probability of intercept
RBW resolution bandwidth
RED Radio Equipment Directive
RFSoC radio-frequency system on chip
RTSA real-time spectrum analyzer
SDR software-defined radio
SFDR spurious-free dynamic range
SigMF Signal Metadata Format
SINAD signal to noise and distortion ratio
TCXO temperature-compensated crystal oscillator
VBW video bandwidth

Sources cited in this handbook

- Analog Devices, **AD9361 RF Agile Transceiver** product data. Receive local oscillator 70 MHz to 6.0 GHz, transmit 47 MHz to 6.0 GHz, 2 by 2 channels. *analog.com/en/products/ad9361.html*

- Analog Devices, **AD9364** product data. The 1 by 1 variant, sharing a register map and package with the AD9361. *analog.com/en/products/ad9364.html*
- **DIFI Consortium**, formed under the IEEE Industry Standards and Technology Organization to advance interoperability between space and ground systems and to prevent vendor lock-in. *dificonsortium.org*
- **IEEE-ISTO Std 4900-2021**, Digital IF Interoperability Standard: a constrained profile of VITA 49.2 that defines what is required rather than what is possible. Tutorial material at *github.com/DIFI-Consortium/DIFI-Certification*
- **ANSI/VITA 49**, the VITA Radio Transport standard for carrying digitized RF with its metadata.
- **GNU Radio project documentation**, for the framework and hardware ecosystem described in Part III. *wiki.gnuradio.org*
- **SoapySDR**, the vendor-neutral hardware abstraction layer named in Part III. *github.com/pothosware/SoapySDR*
- **Signal Metadata Format (SigMF)**, the sample-file plus JSON-sidecar convention recommended in section 40. *github.com/sigmf/SigMF*
- **AMD Zynq UltraScale+ RFSoC and Versal RF series** product data, for the direct-sampling converter rates, the 7.125 GHz input range and the December 2024 announcement quoted in section 6. *amd.com*
- **3GPP Release 17** for non-terrestrial networks, and the **FCC supplemental coverage from space** framework of 2024, for section 61. *3gpp.org* and *fcc.gov*
- **O-RAN Alliance** open fronthaul specifications, for the disaggregated radio access network discussed in section 61. *o-ran.org*
- **Federal Communications Commission**, Report and Order of 22 July 2026 on the upper C-band, for the 3.98 to 4.14 GHz reallocation in section 59. *fcc.gov*
- **Resolution 256 (WRC-23)** and **WRC-27 agenda item 1.7**, for the upper mid-band study bands in section 60. *itu.int*
- **Directive 2014/53/EU** (radio equipment), **Directive 2014/30/EU** (electromagnetic compatibility) and **Regulation (EU) 2024/2847** (the Cyber Resilience Act), for the compliance table. *eur-lex.europa.eu*
- **Section 889, FY2019 National Defense Authorization Act**, and the implementing Federal Acquisition Regulation rules, for the provenance section. *acquisition.gov*
- Berkeley Nucleonics, **ICX-FieldHawk handheld, rugged and USB datasheets**, for every ICX figure in this handbook.

Companion documents

- **BNC-HB-001**, SDRs in the real-time spectrum analyzer field. The deep treatment of SoapySDR, GNU Radio flowgraphs and the calibration boundary that Part III of this handbook summarizes.
- **BNC-AN-102**, wide-area spectrum monitoring, for the intercept arithmetic behind section 32.
- **BNC-AN-105**, RF record and playback, for the two-instrument split named in section 45.
- **BNC-AN-110**, designing a receiver into your own product, for the timing and geometry arithmetic behind Part II.

67 What to take away

Six things, if the rest is forgotten.

- **The market clusters because it shares silicon.** One integrated transceiver family covers 70 MHz to 6.0 GHz and underlies a large share of the products on sale, which is why so many unrelated devices stop at the same frequency and carry the same three impairments. A specification shared by a dozen products is not a differentiator.
- **Sample rate is not bandwidth, and neither is tuning range.** Four different numbers get quoted as a span, and only the narrowest of them governs a measurement.
- **The bit count is a ceiling, not a performance figure.** Effective bits after noise and distortion are always fewer, and the gap between a consumer converter and an instrument one is the largest technical difference in the landscape and the least visible on a datasheet.
- **A reading is not a measurement.** The dividing line in this market is whether an absolute amplitude bound exists, and it is a property of hardware and its traceability chain that no software layer above can supply.
- **A guarantee you can evaluate beats a claim you cannot.** A published intercept relation lets a buyer price a configuration the vendor never considered; a marketing phrase about real-time performance does not.
- **Ask about relationships, not just parts.** Between channels, between units, between the calibration and the samples. These are consistently the least published fields and the most consequential ones.

AND ONE QUESTION TO CARRY INTO ANY EVALUATION

Is device calibration applied before the samples cross your interface, and what bound do you hold to on the samples I receive? Send it to every supplier, including us. The shape of the answers will sort the market faster than any table in this handbook.

VERIFICATION NOTE

The following values in this paper are not yet confirmed against a published Berkeley Nucleonics datasheet and are marked verify in the text. They must be confirmed before this paper is released.

- **That device calibration is applied before the samples reach an open framework is demonstrated, not specified.** It rests on a supplied integration guide with worked examples and result captures, not on a Berkeley Nucleonics specification table. It is the load-bearing claim of Part IV, it is stated here as demonstrated throughout, and a reader who needs it contractually should ask for it in writing.
- **One published column does not close to the digit, and is flagged rather than smoothed.** The Blackman-Nuttall real-time resolution-bandwidth endpoints correspond to about 2.00 times the bin spacing, while the real-time display reports 1.9761 times it. The flat-top column agrees with the same power-of-two grid to better than 0.2 percent, so the grid is taken as exact and the Blackman-Nuttall endpoints as rounded. Awaiting confirmation of which figure that datasheet column was computed with.
- **External figures in this handbook are cited, not measured by us.** The integrated transceiver frequency ranges and the direct-sampling converter figures are the manufacturers' published product data; the standards material is from the DIFI Consortium and IEEE-ISTO. Where this handbook draws a conclusion from them, such as the observation that the market clusters at a 6 GHz ceiling, the conclusion is ours and the reader can check it against the cited ranges.
- **No competitor comparison appears in this handbook, and none is implied.** Device classes are described generically. Where a specific product's specification is quoted it is quoted from that vendor's own published data as public fact, never as a comparison against a Berkeley Nucleonics product.
- **The ICX-FieldHawk does not publish VITA 49 or DIFI support.** The claim appears on a preliminary portfolio page only and is not made here. Section 23 states this plainly rather than omitting the subject.
- **That a routine written against one form factor runs unmodified against another is inferred, not specified.** The handheld, rugged and USB datasheets publish the same programming interfaces, the same SCPI support and the same measurement-function list; the portability conclusion follows from that agreement rather than from a published compatibility statement.
- **Only the 9.5 GHz and 40 GHz tiers carry published specification tables.** Intermediate frequency tiers appear in some marketing material; nothing in this handbook is written against one.
- **Tracking generator hardware, millimeter-wave extenders, onboard AI compute and any EMC measurement function are not published** in any ICX specification table and are not claimed here.

Berkeley Nucleonics Corp.
Test, Measurement and Nuclear Instrumentation since 1963
For details: info@berkeleynucleonics.com or 800-234-7858

BNC-HB-002 Rev A
August 2026