

The Evolution of Interoperability



Making the Dream a Reality.

Anyone who thinks that a standard is synonymous with interoperability is sadly mistaken. The difference is an investment of around \$10 million and five years of hard work. How many standards organisations can afford that?

Nick Hunn

WiFore Consulting

nick@wifore.com

+44 7768 890 148



This work is licensed under a [Creative Commons](https://creativecommons.org/licenses/by-nc-sa/4.0/) Licence. This allows you to copy, distribute and display the contents of this paper, with the exception of the images, or make derivative works as long as the original author is credited.

The Evolution of Interoperability. Making the Dream a Reality.

I've been attending a lot of Smart Energy meetings recently and listening to industry experts talking about the need for interoperability in the connections between smart meters and appliances around the home. I've also been hearing a number of standards organisations trying to promote the message that the concepts of interoperability and a standard are synonymous. That's a very dangerous message, because the two are only very loosely related. Just because you have a standard, it does not mean that products which use it are, or will become interoperable.

To understand why equating a standard with interoperability is a fallacy, let's start with an analogy. In many ways, a standard is like a language. So we could define English, or French or Russian as standards. The standards bodies would then claim that everyone who speaks the same language is interoperable. I'd disagree. The language defines the grammar and the vocabulary, but you only have to listen to a Democrat and Republican senator debating health reforms to understand that speaking the same language does nothing to promote interoperability. If anything, a standard provides the tools to ensure that conflict is more, rather than less likely to occur.

Interoperability is about working together seamlessly. To achieve that requires more than just a standard. It needs a set of interoperability tests and the testing tools to confirm compliance with those tests. These don't generally come with a standard – they need to be put in place to support it. That entails time and money, which means most standards can't support them until they're already fairly well established. Industries like Smart Energy demand interoperability, as they want the meters they install today to work with devices that customers may install in ten or twenty years' time. But if they want to achieve it, they, need to understand how this process works.

Let's start with a definition: Interoperability in this context is the ability for any two devices to connect together and support a meaningful dialogue, regardless of who made them, how they're used, or what else they may have connected to in the past. It is the aspiration to that ideal state where there will be a flawless interaction between products created by different manufacturers.

That's not an easy definition to visualise. Interoperability can be a difficult concept to grasp, especially for a wireless connection. The primary reason for that is although wireless standards basically replace a cable, at the same time they offer the promise that once "connected", widely different devices should be able to "communicate" with each other. I use quote marks for both of those terms, as when you think about them, they're far from obvious. If you connect a weighing scale to a mobile phone, what should it do? What does a connection between a smart meter and a washing machine mean? Or one between a pair of running shoes and a headset? At the highest level, interoperability has to rely on

an application, which goes beyond the scope of wireless standards. However, interoperability has become a desirable tick box in many industries, regardless of whether it makes any sense.

To explore what wireless interoperability means, let's start with the simplest wireless scenario of a broadcast service. A good example is a television, where programs are transmitted in a defined format by a few large broadcasters to a range of passive receiving devices called TVs. The way in which the programs are broadcast is determined by a number of different standards – typically NTSC, PAL or SECAM. Because broadcasts are normally confined to national boundaries, there's not an obvious interoperability problem. However, those different standards meant that TV manufacturers used to have to make different, incompatible models for different countries. After thirty years of relative stability, newer HD formats have arrived which are not compatible – if you have an old TV you can't display them. Nor is there any way you can upgrade your TV to an HD format. Which highlights two points: which are that both ends of the wireless link need to support the same standard if you're going to get interoperability, and that interoperability is not the same as, nor does it confer backwards compatibility.

For broadcast TV, interoperability is relatively easy, as a very few companies control the broadcasts, and as soon as you turn on your TV it receives the signal. It does not need to connect, or pair, or negotiate encryption keys. It's an embedded device that has been designed to do one thing. The next level of complexity is where devices can have a two-way conversation, transmitting as well as receiving. That means they need to connect and establish a link between each other.

An everyday example of this is a mobile phone and the cellular network. This is a simpler interoperability problem than it might first appear, because phones do not connect directly to each other. Instead each connects with the network infrastructure, which transfers the data between the two (or more) appropriate handsets. Handsets and base stations have to conform to industry standards. For most phones those are defined by ETSI – the European Telecoms Standards Institute, which is responsible for the GSM and 3G standards which are widely used. It's not the only standard, but it accounts for over 4 billion connections throughout the world.

Despite those numbers, there are not that many different mobile phones, and even fewer different base stations that they connect to, because the industry is controlled by a relatively small number of companies. One of the reasons for the small number of companies is the cost of implementing the standards, plus the cost of testing them. Before a mobile phone is brought to market it needs to pass a stringent set of qualification tests, which can cost anything up to \$1,000,000. At that point it becomes legal to sell it. However, before a network operator will sell it to you, they insist on it passing a further set of interoperability tests. In Europe that's the [GCF test](#) schedule; in the US it's the [PTRCB](#), both of which are of similar complexity. These extend the normal tests to check that the phones correctly implement the features that are critical for

working on real networks. Both of these provide significant barriers to entry for new companies, but between them they provide a high level of assurance for users and operators that phones will work. That's very important for network operators, because if phones don't work, then they won't make any revenue from calls. Yet despite that, most users have experienced calls failing to connect, or had issues in sending pictures between different networks. And that's despite the fact that there's a limited number of players, that phones can only connect to a limited number of different base stations and that every product that gets to market is extensively and expensively tested.

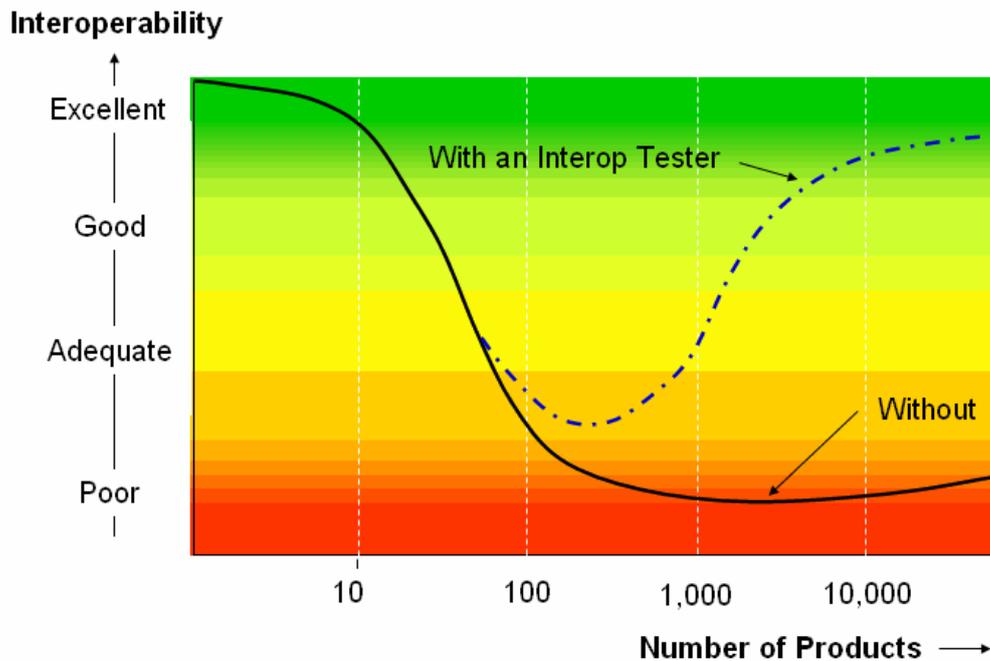
Which brings us on to short range wireless. It doesn't matter which standard it is – they all have the same issues. Users expect a wide number of products to be able to connect to each other, for them to work out what they can do and then share information with each other, whether that's a piece of data like room temperature, a voice stream, a data file or a video. Interoperability is the expectation. Regardless of who makes an individual device, users believe that it should connect to any other device which supports the same wireless standard, within the bounds of what is reasonable expectation (and I still don't expect a pair of running shoes to talk to my Bluetooth headset, but one day someone will).

Specification writers try hard to make sure that their standard supports interoperability. That means they need to define how devices connect, what protocols they use to transfer data between each other, how that data is formatted and how to recover if things go wrong. Defining that for a wide range of applications is difficult, which is why most wireless specifications are several thousand pages thick. The problem is that this level of complexity invariably means that different implementers take slightly different approaches, and because they're human, make slightly different mistakes. Out of 2,000 pages, it only needs two developers to interpret one line of that specification differently and you have an interoperability problem. That's the main reason that devices don't always work with each other in the way we expect.

That's the reason we have problem. What is counterintuitive about interoperability is that it changes during the life of a standard. In the early days of a standard there are very few products on the market. Generally they're being produced by the companies that have dedicated resources to writing the standards, so they're keen to make sure that these first products are interoperable. (That's just good vested interest. If they're not, the negative publicity doesn't attract other companies to use the standard, and consumers are put off buying it.) To ensure interoperability, these manufacturers will test their products with each other, making any modifications necessary to demonstrate interoperability. The important point to note here is that these modifications may not make the products comply with the standard. They are fixes to solve problems, and as such may perpetuate deviations from the core standard. These will probably be ignored if everything works, as without a test system that checks compliance to a standard, expedience tends to win.

As the number of products starts to grow, the task of testing everything against everything else becomes impossible, as it grows factorially. Major manufacturers will still perform extensive testing of their flagship products, but in general interoperability starts to take a nosedive. The other thing that happens is that as more and more manufacturers start to write

The Evolution of Interoperability



protocol stacks and profiles, each tends to deviate slightly from the standard, because of minor differences in interpretation and implementation. Rather than testing these rigorously against the specification, effort tends to be put into ensuring interoperability with what each manufacturer sees as the market leading product. That results in more patches to make their stack work. If in turn they become successful, other manufacturers will do the same thing against that product, running the risk that de facto implementations start to diverge from the specification.

As more and more products and manufacturers come to market, this leads to sets of products that work with others in their group, but which exhibit a steadily decreasing level of general interoperability against the standard.

Many standards attempt to address this by designating a small number of "golden units". These are products that are meant to be the best case implementations, and which act as a reference against which other products should be tested. There's generally strong competition between manufacturers for their products to be chosen as golden units. That can be counter-productive, as it results in a rush to generate products early in the life of the specification, when they are likely to have had very limited testing.

It also leads to another problem, which is the issue of waivers. If a product fails to interoperate with a golden unit, but the manufacturer

believes that their product conforms to the standard, they may ask for a waiver. In theory this should lead to the golden unit being updated. Far too often it results in a growing grey area around the golden units, where once again market practice diverges from the word and intent of the specification. It's not uncommon with the golden unit approach to find two products that comply with the standard, but which cannot talk to each other, or two products that are interoperable, but where neither conforms to the standard. There have also been examples of two golden units from different manufacturers that are incapable of working with each other, but which are still considered acceptable for testing.

The solution is for a standards organisation to invest in automated interoperability testing equipment which is designed to test every product against the specification. It's not an easy option to choose. The Bluetooth SIG made that choice in 2005, when they saw a growing number of interoperability issues. The result was their Profile Tuning Suite (PTS) – an automated piece of testing software that runs on Windows, and which can be used to test (and help develop) new Bluetooth products. Because it tests products directly against the specification there should be no question of deviation. If a manufacturer thinks there is a problem, then that can be checked against the standard, and the PTS updated to correct an error, or an errata can be raised to clarify the specification itself. It provides a closed loop system that helps to ensure that all new products move closer to the standard, rather than drifting away.

It's not a trivial decision to develop this type of software. Over five years, it has cost the Bluetooth SIG around \$10 million to develop it to the level at which it operates today. That cost doesn't diminish with time, as there is an ongoing expense to maintain and update it, as well as extending it to cover new releases of the Bluetooth specifications. Nor is \$10 million the whole cost to date - member companies and test houses have probably contributed an even greater amount in terms of engineering resource that they have provided. But the results are dramatic. Since it's been in place, it has been used to test over 8,000 products and has resulted in a significant improvement in interoperability, as indicated by the diagram.

No other short range wireless standard has even started down this route. That means that Bluetooth (and Bluetooth low energy, which uses the same tester) have a five year lead over other short range standards like ZigBee and Wi-Fi in terms of being able to guarantee interoperability.

When initiatives like Smart Energy talk about interoperability they need to understand this distinction. A standard does not confer interoperability. In a mass market, interoperability can only be achieved by stringent testing against the specification itself. For that you need a test system, backed up with an enforcement policy that can remove non-compliant products from the market. Unless an organisation has already embarked upon putting those into place, it will take them around five years before they can even start to claim that they have a handle on interoperability. Interoperability does not have short cuts.

In conclusion, if you need interoperability, make sure you're asking the right question. Slick PR from a standards body is no substitute for doing it properly. If a standard tells you it's interoperable, don't believe them until they show you their interoperability test system and enforcement program. If they can't, it's likely that they'll follow the curve of increasing interoperability problems and growing customer disenchantment.

Nick Hunn

Dec 2010