

So you want to create a standard?

Professor William Webb, CEO Weightless SIG

Introduction

Perhaps you have a vision for how a particular market will evolve. Or perhaps you have spotted a missing element in an existing market. Even better you know how to deliver into this vision using a new approach or a new technology. The market must be sufficiently large to justify a new technology and the technical approach must be implementable and cost effective. Assuming these criteria are met, should you go it alone, or should you create a new standard?

Standards tend to result in larger markets, but these are shared among multiple players. The key question is then whether all of the “smaller pie” of going it alone is worth more than a slice of the “larger pie” of taking a part in a standard. This is a very difficult call to make as predicting the outcome under either scenario will be subject to huge uncertainty. However, in some areas such as wireless links, standards are virtually mandatory – indeed, all of the successful wireless technologies in use today are open standards (eg 3G, LTE, Wi-Fi, Bluetooth, DECT, etc).

Having made the decision to move ahead, there are a number of choices to be made including the IPR policy, whether the standard will start with a clean sheet of paper or a pre-designed solution and which standards body to use as the vehicle. Underlying all these decisions is the need to get buy-in from others if the standard is going to succeed. Each of these key decisions is addressed next before returning to the issue of getting enough buy-in.

The tension between sharing and owning

Having made the decision to move to a standard you have decided to give up part of the pie in return for a larger overall pie. However, the risk is that the share of the pie becomes smaller than expected due to more intense competition. One way to insure against this, as well as to gain some useful additional revenue, is to hold key IPR in the form of patents.

IPR and standards is a very fraught mix. Everyone wants to hold key IPR and nobody wants to pay royalties to anyone else. If this becomes unbalanced and others involved in the standard have to pay large royalty fees to one of their competitors they may choose to leave and even establish a competing standard. If there is dispute this can embroil the standard in lengthy legal proceedings which are costly, delay production and ultimately can persuade others that the standard is more trouble than it is worth.

Another problem with IPR is that it can lead to a poor standard. If major contributors are all pushing technological choices for which they have IPR then the tendency is for all the conflicting ideas to end up in the standard as “options”. This leads to ill-focussed standards that are difficult to implement and generally sub-optimal.

One way around all of this is to declare the standard “royalty free”. This is a route followed by Bluetooth. Essentially all those involved in the standard agree not to charge royalties to each other on activities related to the standard. This avoids the downsides but can seem like a harsh approach for the company who started the work in the first place, probably has the most IPR and is looking for some recompense. In the case of Bluetooth this company was Ericsson, but they were content to waive royalties on the basis that if Bluetooth was successful they would sell more cellular handsets which would be of greater benefit than IPR fees from Bluetooth (whether this outcome actual transpired is far from clear!).

A key determinant, discussed further below, is achieving buy-in. Offering up a near – complete standard which is royalty-free can help encourage others to join since it avoids the IPR penalty often associated with later entrants to any standard. Hence, if buy-in is critical this may be a necessary option.

Blank sheet of paper or injected “v0.6”?

A classical approach to standards would be for someone to determine a need for a new technology – for example to deliver M2M services – and commission a standards body to develop a standard to suit. This sometimes happens, for example, with the European Commission noting a need and asking ETSI to develop a standard. In principle, ETSI starts with the proverbial “blank sheet of paper” and those interested in being involved design the system from scratch.

In practice it rarely works this way for two reasons. Firstly, the need for solutions rarely arises in isolation of ideas as to how to deliver them. Secondly, the blank sheet of paper approach is often ineffective. We discuss both below.

The understanding of problems and the technical solutions tend, in most people’s minds, to be related. A problem with no known practical solution (eg a robot that can do the ironing) is not ripe for standardisation. A new area of communications to which existing communications technology can be applied is better suited. As a problem, or need, is studied, solutions will be considered and the need and the solution iterated towards a viable outcome. This may only be an outline, but nevertheless, further advanced than the “blank sheet”.

The blank sheet is also very slow. The hardest decisions for a standards body to take are the bigger ones – such as should a communications system be based on single or multiple antenna solutions. The easier decisions are the details - such as should this timer be 10s or 30s. This is not only because the bigger decisions need more study but also because generally they will have a major impact on the IPR drawn into the standard and may favour one company’s expertise over another. Blank sheet solutions tend to result in a number of competing systems proposals being put forward by participating companies and then a selection process that tends to end up compromising by adding in parts of many of the solutions. Sometimes, as with ultra wideband, the different proposed solutions are so far apart compromise is not possible and deadlock is reached.

This all suggests that starting with a partially complete standard is more likely to succeed. The experience of Bluetooth which started with a “version 0.6” standard that

was around 50% complete is useful case study. Firstly, any partially complete standard needs to be complete enough to get over the initial system-level decisions but still leave plenty of scope for other participants to influence and enhance the standard. Secondly, and most critically, it needs to be clear that the company delivering the partial standard is not gaining disproportionately from the IPR they might hold. In the case of Bluetooth and Weightless both leading companies decided to offer a royalty-free licence to the technology, overcoming this concern as discussed above.

Which standards body – or invent a new one?

A tactical decision to be made is which standards body to use. There are broadly three categories:

- *Formally recognised standards making entities.* These are bodies such as ETSI that have some specific rights, often given by national or international Governmental entities, to make standards. Once made, their standards automatically apply and are recognised across the relevant geography.
- *Informally recognised entities.* These are bodies such as the IEEE that have no formal remit but because of their previous success in delivering standards, are generally recognised as standards bodies.
- *Bespoke entities.* These are bodies set up specifically to deliver a particular standard. The Bluetooth SIG is a good example.

As always, there are advantages and disadvantages with all these approaches. Broadly bespoke entities are less bureaucratic and hence faster, better able to cover associated areas such as certification and marketing, but bring less credibility and require substantial effort to set up.

With an existing standard that is to be enhanced it is clear that the relevant standards body should be used. For example, the IEEE802 committee is the only place where enhancements to the 802.11 Wi-Fi standard could be made. Similarly, the 3GPP committee is the natural place to deliver 3G and 4G-related cellular standards. However, recognised bodies have recently been less successful with new ideas. ETSI has failed to deliver much in the way of new standards in the last decade and the IEEE has seen much conflict and delay around new standards.

It is not always an “all or nothing” decision. A standard developed in a bespoke body can subsequently be moved to a recognised body where it may be ratified, gaining the benefits of legal and regulatory recognition. In recent years, recognised bodies have also offered “incubators” for bespoke standards, allowing them to use the resource of the recognised body while acting outside of the bureaucracy, with a defined path to moving back into the standard approach should that be desired.

Ultimately, it tends to come down to the need for speed. Existing bodies, however hard they try, tend to be slow. Any new standard is attended by many members some of whom may wish to delay, stall or change the direction of the concept. The rules of the club often require consensus or at least majority voting rather than dictatorial decision. If the idea for the standard has come from a large organisation there may be less need for haste as the standard will be just another product line and those involved

part of an R&D or standards team. However, if the idea is from a start-up then the duration of available funding will be short and the need to deliver rapidly critical to its survival. Of course, this is simplistic and there may be many other reasons why timing is more or less important.

Pulling off a bespoke standards body is not to be undertaken lightly. It needs to rapidly establish sufficient credibility that it becomes seen as a standard and it needs a framework to operate within covering legal, website, IT, membership, finance and much more. Credibility is a particular issue, and is discussed in the next section. In essence, at launch the new standard needs to have sufficient committed companies to be seen to be important and that list needs to grow with sufficient rapidity over time that the standard is considered to have “momentum”.

The framework for the body would merit a detailed paper in its own right. A summary is provided below.

Legal framework. The framework for a standards body is complex. Legal documents need to cover the IPR policy, governance, role of the Board members, conduct of the other members (often encompassed in bylaws), confidentiality, copyright and more. Most of these will be scrutinised in detail by the lawyers of the prospective members and there will be endless discussion and iteration. Although existing documents, such as bylaws from other standards bodies, can be used as templates, they always seem to need modification. Experience suggests that establishing a legal framework will take at least a year and likely cost in the region of half a million dollars. However, with some clever drafting of interim agreements it is possible for work on the standard to get underway before all the legal documents are fully agreed.

Website. Naturally a standards body needs a website. This plays a number of roles. It is a recruitment vehicle for new members, a set of resources for existing members and a portal to secure document storage. Establishing and maintaining it to a level that will provide credibility will take months and need a range of resource, all costing hundreds of thousands of dollars.

IT. A standards body needs good collaborative tools. It needs to share documents securely, allow for editing, good version control, provide email servers, maintain membership lists of sub-groups, deliver conference phone facilities and more. While much can be done with tools like Google Documents, there is still substantial administrative work required.

Marketing and PR. A standard needs to be seen to be growing, typically measured in terms of membership numbers. This is discussed more below. It cannot happen without excellent marketing to “spread the word” and evangelise the new technology. The importance of getting this right cannot be overstated.

General administration. There is much to do. Fees need to be collected, members updated or removed, enquiries answered, events arranged. At least one full-time employee is likely to be needed.

Creating the standard. Amongst all of this, it is easy to forget that the fundamental reason the body exists is to deliver a technical standard. This needs careful

management, work needs to be segregated into sub-groups, chairs and key members need to be “persuaded”, conformity needs to be maintained and much more. Above all, if the advantages of speed are to be delivered, pressure needs to be carefully applied, key decisions made quickly and some more complex features may need to be deferred to subsequent updates. All of this really needs a “Chief Architect” who informally, but through their personal credibility, can cajole, break deadlock, set targets and herd disparate volunteers towards a common goal all buy into. It is a rare individual that has the technical skills, personal credibility, management expertise and inter-personal touch to bring this off.

Achieving sufficient buy-in

We now turn to a topic touched upon earlier – that of credibility. A standard is only a standard if enough people believe it to be. It is only successful if enough companies engage and make products or deliver solutions. There are many more failed standards than successful ones, and most fail not because of technical shortcomings but due to a lack of market traction.

To be successful a standard must, at its heart, be a *shared vision*. If enough of the key actors *believe* that a developing standard fits their interests by opening doors for profitable operation then they will act in such a way as to make the standard a success. New standards tend to target new markets, so there is no sure-fire way of knowing that a new standard will deliver profitability. Key players need to have *faith* based upon their experience and expectation. It is easier to keep the faith when surrounded by fellow believers. Perhaps a standard is closer to a religion than it is to a user manual. For any new standard the key question is whether it can build a new religion with enough believers to create a virtuous spiral of ever-growing support, increasing equipment availability, ever falling costs and a stream of innovative new applications.

Ideally a new standard needs a few key “anchor tenants”. These are companies that add credibility through their size or previous history – companies like Google or Cisco. At the start, having a relatively small number of total members can be helpful as this allows the work on the standard to progress more rapidly. Then, as the standard is complete and the focus turns to implementation, a growing number of members is needed both to undertake the implementation work and demonstrate that the standard is “going somewhere”. This generates the virtuous spiral that leads eventually to a successful standard becoming globally dominant in its particular space. Again, Bluetooth did all of this particularly well.

Making this work, from inception, to finding the anchor tenants, to growing the membership, is all marketing of one form or another. Initially it may be a few key proponents using their contact base to persuade individuals in key target organisations. As time goes on it might become more about classic PR channels using journalists, websites, conferences and so on. While the channels may be familiar, the marketing techniques need to reflect the fact that there is no product here – this is about selling a vision and the reason for a company to be a part of it. Unless the standards body is well endowed financially, all of this will need to be done on a very tight budget, possibly without any dedicated resource.

Fundamentally it needs the right people. Since there is no product or service to sell then it comes down to the persuasiveness of the proponents. These people need personal credibility to open doors, excellent communications skills to get the message across in the most persuasive manner, enough drive and energy to give variants of the same message repeatedly over months and years, excellent connections and a deep understanding of the technology. They are very rare breeds indeed, and unless the standards team has at least one on board at the start it will likely not succeed.

Coupled to this is a marketing team that can make links to journalists and analysts, handle all the various social media channels, maintain the website and ensure it is search-engine optimised, deliver videos, magazines, member profiles and other content material, set up events such as press launches, translate material into other languages and so on. In the early stages they need to generate enough material to give an impression of growth at a time when little other than standards development is taking place.

Complicating all of this is that the marketing must be aligned with the corporate needs of the key players in the standards body who will likely have differing agendas.

Achieving credibility and momentum is the most critical element for any standard, and yet the hardest to define, deliver and measure.

Will it succeed?

So you have spotted a need for a new standard and devised a technology that can deliver against that need. You have determined an IPR policy that steers an effective middle-ground between owning too much and encouraging others to join. You have delivered a semi-complete standard and encouraged others to work on completion in a collaborative, non-competitive and time-effective manner. You have determined which standards body to use, or perhaps established your own. Above all you have promoted the concepts to hundreds if not thousands, through every media channel available. You have been busy!

Will you succeed? The simple answer is probably not. Most standards fail, predominantly because they fail to gain credibility. Some fail because of competition – for example WiMax ended up competing with LTE. Some fail due to delays in standardisation. A few fail due to poor technical design. But with enough credibility and support, a clear vision, a strong need and good technical solution you might just change the world.

Appendix – The choices made in Weightless

The Weightless standard was developed between 2011 and 2013. This appendix provides a case study of the decisions made within the framework set out in this paper.

Weightless was first conceived by Neul who perceived there to be a lack of a long-range machine communications technology able to deliver on the vision of machine-to-machine (M2M) or the Internet of Things (IoT). With all the early work on Weightless performed by Neul, one route would have been for it to remain a proprietary technology. However, experience shows that there are no successful proprietary wireless technologies – all of the wireless systems we use routinely are open standards. This made it clear that Weightless needed to be an open standard in order to deliver on its vision.

The Neul team then considered the most appropriate IPR position. The design work done to that point suggested that there would be limited prior art because many of the techniques used were over 20 years old and so any relevant patents would have expired. Of those areas where there was IPR, much of it was patented by Neul. Machine communications systems need to be very low-cost (typically less than \$2) and hence there was little room for royalty payments while remaining within this limit. The Neul team had also been involved in Bluetooth where some of the success of the standard was seen to be due to the royalty-free regime adopted. The team considered whether royalty-free would encourage others to join, and decided on balance that given most of those joining would not have relevant IPR in this space that such a regime should be attractive. Based on these considerations, eventually a hybrid approach was selected. Terminals would be royalty free, however base stations would be licensed on a classic “fair, reasonable and non-discriminatory” (FRAND) basis as widely adopted in bodies such as ETSI.

The clean sheet of paper versus “v0.6” argument was, by this stage, not relevant. Neul’s initial work to show that there was a possible technical solution for this market had resulted in a v0.6 specification already. It made little sense not to use this as the starting point. Hence, Neul provided a specification that was considered around 60% complete¹. This approach was made simpler by the IPR policy, meaning that those joining the standards body would not have to pay royalties to Neul, nor would they be disadvantaged by the standard not including their own IPR as they would not get any royalty payments on this in any case.

Perhaps the most difficult decision was whether to use an existing body (the most likely would have been ETSI or the IEEE) or to create a new standards entity. With the team’s experience of Bluetooth, which had created its own body (the Bluetooth SIG) there was some natural bias in that direction. Another key determinant was speed – Neul was funded by VCs and the time horizon to deliver working equipment conforming to an open standard was short. Based on this, Neul decided to attempt to establish a new body called the Weightless SIG, but to keep this decision under

¹ Determining completeness of standards is rather arbitrary. Time has shown that the version 0.6 was only about 25% complete in terms of page count, nearly 100% complete in terms of system design and perhaps around 50% complete in terms of details such as message structure.

review. Even now, some two years on, there are periodic reassessments as to whether the standard should be moved into existing bodies such as ETSI. However, as time goes on and the Weightless membership grows this becomes increasingly unlikely.

Setting up the Weightless SIG has been a long and arduous process. The first step was to recruit the “launch Promoters”. These were credible companies that would form the initial Board and show, at launch, that the standard had support outside of Neul. These companies could not be found through marketing as, by definition, this was pre-launch, but had to be recruited through personal contacts, meetings and detailed discussion. This stage is really tough, with key companies opting in and out through the process (often as individuals were moved internally) and legal teams working on initial “memorandum of understanding” documents that would form the basis for subsequent legal work. This stage took about 6-8 months at the end of which ARM, CSR, C&WW and Neul were on board (actually, at this stage two other companies were also involved but subsequently withdrew predominantly because of inability to agree legal terms).

There then followed a year-long process to agree the legal framework and IPR policy for the SIG. That this much time (and associated cost!) was needed might seem incredible, but every legal concern had to be addressed and circulated to all Board members, who typically took a month to respond, often with subsequent concerns. Changes of legal teams in some cases did not help. Thankfully, the SIG had the foresight to put in place an “Early Adopters Agreement” which was a temporary legal structure that allowed work on the standard to commence before all the formal legal documentation was in place. This allowed the standards body to get going around February 2012 even though the full legal framework did not get signed off until November 2012 and put in place by February 2013. Indeed, the decision to set up a temporary framework, although incurring additional cost, was with hindsight one of the most important ones made in ensuring the standard did not stall at an early stage.

With the Board in place, members needed to be recruited in sufficient numbers to populate the various sub-groups and provide enough effort to draft the standard. This was achieved through a launch event run by a Cambridge organisation called Cambridge Wireless in September 2011. This was extremely helpful as they could market the event and manage the logistics. With their reach over 100 people attended, and from this the first 50 members of the SIG were recruited. Using an existing organisation was another key decision in getting the SIG off the ground.

Events and legal negotiations cost money, and an issue for the SIG over its entire existence to date is financial. Indeed, any successful SIG has typically a couple of years of debt followed by many years of surplus. In early years expenses are multiple – legal, events, marketing, website, IT, etc. However, revenues are limited to the relatively few member fees. In later years costs are less and revenues grow quickly with more members and the possibility of revenue from certification, events, sponsorship and other areas. The only viable solution was a loan from the Board members – in this case from Neul. Without this, Weightless would not have been able to launch.

From February 2012 work started in earnest on the standard, with 6 sub-groups, each meeting weekly to progress their area. Quarterly “Plenary” meetings provided a

regular cycle of updates to the baseline specification as well as injecting the urgency of deadlines into the drafting process.

By the end of 2012 it was clear that the standards work was near-complete and it was time to grow the membership and the credibility more rapidly. That meant marketing. But the SIG was stuck in a chicken-and-egg situation where marketing to get new members required money but the money to pay for the marketing would only come after the new members joined. This was resolved through an innovative marketing agreement with a skilled team who agreed to be paid through a percentage of member fees rather than against hours worked. Although still in its early days, this arrangement looks extremely promising and may well be another of those key decisions that changed the future of the SIG.

The SIG has needed much administration, IT support and so on – along the lines of that discussed in the paper. It also needed a CEO. This was both for legal reasons, to sign key documents, etc; but also for credibility and marketing purposes with a key figurehead representing the SIG in major venues and being seen to provide leadership.

Much more could be written about the story of Weightless but this account sets out the key strategy and decisions.

As mentioned in the article, a key element is achieving sufficient buy-in from others to make the standard credible. This is an on-going activity within Weightless. We have learnt that broadly, there is a need to ratchet upwards in terms of company size. Initially, the large companies see the activity as something to observe. Smaller players are happy to get involved and some mid-size companies may be interested if it closely fits their strategic goals. As the mid-size companies join the larger ones start to take increased interest and it may be that one of two of the larger companies can be persuaded to join because of competitive concerns about the mid-size entities. As these larger companies come on board, other larger players perceive a standard that is going somewhere and start to join themselves. This is naturally a slow process requiring marketing, lobbying, pursuit of individual companies and much direct discussion with key individuals. It has been helped by the “anchor tenants” that form the Weightless Board, including ARM and Cable & Wireless. As the standard is completed and key milestones such as silicon availability are met, the process becomes easier both because marketing can be more specific and because larger players can see that clear progress is being made. Building credibility further is now the top priority for Weightless.