DAVID RUDIN

Coffee & Standards

A Practitioner's Guide to Standards, IP, Governance, and the Long Game

FIRST EDITION

Coffee & Standards

A Practitioner's Guide to Standards, IP, Governance, and the Long Game

DAVID RUDIN

First edition, 2026

Disclaimers

Not legal advice. This book is a practitioner's guide to the governance, intellectual property, and negotiation dynamics of technical standards. It is not legal advice and does not create an attorney-client relationship. Readers should consult qualified counsel regarding their specific situations.

Table of Contents

TABLE OF CONTENTS

PART II: PATENT POLICIES — THE CORE OF STANDARDS IP

PART III: GOVERNANCE, PROCESS, AND GETTING THINGS DONE

Chapter 11 — Standards Governance Structures

APPENDICES

SOURCE MATERIAL INDEX

Based on the "Coffee & Standards" training series conducted by David Rudin (Microsoft CELA), October 2024 – July 2025, supplemented by additional source material.

ABOUT THE AUTHOR

David Rudin is an Assistant General Counsel at Microsoft, where he has spent 20+ years focused on standards, open source, and intellectual property. He has represented Microsoft in dozens of standards bodies and open source foundations, and has helped shape the IP policies of many of them.

Foreword

This book is the result of a series of informal training sessions I ran with a small group of attorneys at Microsoft. We called them "Coffee & Standards" — an hour every week or two, walking through the IP, governance, and negotiation dynamics of technical standards. No slides. No formal curriculum. Just a screen share of actual agreements, actual policies, and whatever was top of mind that week.

The goal was simple: get people to the point where they could pick up a standards agreement, read it, and figure out whether something was outside the range of reasonable.

That range, it turns out, is pretty wide. Standards work has matured enough that you don't see much truly crazy stuff anymore. But it still shows up — in disclosure obligations, governance structures, voting mechanics, patent commitments. And when it does, you need to be able to spot it.

I've been doing this work for over two decades. Long enough to have seen most of the patterns more than once. The benefit of that tenure is pattern recognition. I can look at a particular clause and tell you it exists because of a negotiation between IBM and Microsoft fifteen years ago, and that the language percolated from one organization to another until nobody remembered why it was there. These organizations are deeply interconnected that way. Terms travel. Context doesn't.

That's the gap this book tries to fill. Not just the legal analysis — you can get that from a careful read of the agreements themselves. What you can't get from the text is the judgment layer: why a provision exists, what problem it was solving, what happens when you push on it, and what the person across the table is actually worried about.

A few things I've learned that shape everything in this book.

First, the skills that make you effective in litigation or bilateral contract negotiation will actively hurt you in standards. There's no judge. You'll never convince someone to take a position against their own interests through clever argument alone. If you go for the jugular in the morning session, you'll be sitting across from that same person at lunch and again next quarter. These are repeat players in a long game.

What works instead is something closer to diplomacy. Back-channeling. Subtlety. Empathy — genuinely understanding what every participant in the room needs, what they can give on, and what they can't. If you do this well, you'll find other people taking your positions for you. You can sit back and watch it happen.

Second, a lot of what we do is theater. Most standards-related patent commitments will never be litigated. The releases, the pledges, the non-asserts (promises not to enforce patents) — they exist to land properly and manage perception. Sometimes the legally optimal approach isn't the best approach. Sometimes a vehicle that is technically weaker but broadly trusted will accomplish more than a bespoke agreement that's technically airtight but nobody recognizes. The MIT license isn't perfect. Neither is the W3C patent policy. But people know them, and that matters.

Third, process is substance. In standards, how you get to a decision matters as much as what you decide. Due process, right to appeal, consensus — these aren't just governance niceties. They're what gives a standard legitimacy and, not incidentally, antitrust protection. When process breaks down, everything downstream is at risk.

This book is organized to build from the ground up. Part I covers the foundations — what standards are, the organizational landscape, how intellectual property maps onto them, the due process and antitrust framework that makes the whole system legitimate, and the increasingly blurry line between standards and open source. Part II goes deep on patent policies, which is where most of the complexity and risk lives: necessary claims, RAND and royalty-free regimes, exclusions, non-asserts, and the practical pitfalls that recur across policies. Part III covers governance, decision-making, and the mechanics of getting a specification from first draft to publication. Part IV turns to practice and reflection: multi-party negotiation, practical advice for counsel, lessons from two decades of this work, a look at where the field is heading, and some closing thoughts.

Three appendices follow the main text. A glossary of standards terminology. A standards-engagement checklist for counsel. And a case study — the HD DVD vs. Blu-ray format war — that shows how the concepts in the earlier parts play out when the stakes are real.

This is a book for practitioners — people who are starting, advising, or participating in standards engagements. It covers the lifecycle from forming an organization and selecting an IPR policy through developing and finalizing a specification. It is not a book about what happens after implementation, when disputes over royalty rates and licensing terms move into litigation. There is surprisingly little case law on the patent policy terms themselves. Most standards disputes that reach courts are about rates and licensing mechanics — important topics, but beyond the scope of what we're doing here.

I've tried to write it the way I would explain it over coffee. Lead with the point. Give you the tradeoffs. Use real examples. Skip the parts that don't earn their place.

If you're an attorney advising a client on a standards engagement for the first time, start at the beginning. If you've been doing this work and need to sharpen a particular area — exclusion mechanics, RAND rate dynamics, voting structures — jump to the relevant chapter. Most chapters are designed to stand largely on their own, though Part II builds on concepts introduced in Part I, and the case study in Appendix C draws on the full arc.

One last note. The agreements and policies discussed in this book are real, drawn from organizations like the Joint Development Foundation, W3C, OASIS, ISO, and others. Where I've included my own analysis or interpretation, that's what it is — one practitioner's view based on long experience. Reasonable people will disagree on some of this. That's fine. The goal isn't to give you the answer. It's to give you the framework to find your own.

Let's get to work.

David Rudin

Chapter 1 — What Is a Standard and Why Should You Care?

Standards are everywhere. If your kid has a bicycle helmet, flip it over. You'll see certification marks and references to safety standards that define how much impact the foam has to absorb. Your car's crumple zones, the electrical outlet on the wall, the USB port on your laptop — all built to standards. Most of the time, nobody thinks about this. That's the point.

The standards we deal with in the technology space work the same way, just applied to software. They describe how two systems connect and communicate. Think of it as a software equivalent of a power outlet. The standard defines the shape of the plug, the voltage, the amperage. It doesn't care what you plug into it.

1.1 Standards Come in Multiple Flavors

Not all standards are the same, and it helps to separate them into rough categories.

Measurement and construction standards are the oldest category, and they're worth naming first because they shape how everyone else thinks about the word "standard." Weights and measures — the meter, the kilogram (the international prototype kilogram in Paris famously appeared to lose tens of micrograms of mass relative to its sister copies over the course of a century, which is part of what eventually drove the 2019 redefinition of the SI base units in terms of physical constants), the volt, the kilowatt-hour — sit underneath essentially every modern transaction. A pound of coffee means the same thing in Seattle and São Paulo because someone, somewhere, agreed on a definition and built an institutional apparatus to enforce it. Building codes are the close cousin: how thick a load-bearing wall has to be, how far apart studs can sit, what gauge of wire runs through which conduit. These standards rarely come up directly in technology practice, but they're the cultural template — when a non-specialist hears "standard," this is usually what they're picturing. It's worth knowing they're there because, while this book focuses on the lessons drawn from interoperability standards, all of these flavors share the same foundation: an agreed-upon definition, an institutional process for maintaining it, and a community that relies on it.

Health and safety standards are the ones most people think of next. Bicycle helmets. Bumper crash ratings. Factory safety requirements. These carry regulatory weight and often end up embedded in law. We won't spend much time on them, but one thing worth noting: these standards become tools of industrial policy. American automakers push for worldwide adoption of European bumper safety standards because it levels the playing field. If Toyota has to build one bumper design that meets the European requirement, they're competing on even terms everywhere. A local manufacturer in Brazil, say, can sell a cheaper car because they only need to meet a lower national standard. Toyota loses on cost. Harmonized standards fix that.

Management and process standards are the ones you get audited against. ISO 9000 is the canonical example — it's one of the most widely adopted standards in the world. These standards don't tell you exactly what to do. They tell you what capabilities and processes you need to have, and then you describe how you meet them.

Here's the way I explain it. Imagine you have two mountains and you need to move goods from the top of one to the top of the other. The standard says you need a mechanism to do that. If you're a small operation, maybe your plan is: we hike down and hike back up once a month. Fine. Maybe you built a bridge. Maybe you use a helicopter. The standard doesn't care how — it just requires that you have a plan, you can describe it, and someone can evaluate whether it's reasonable.

The one closest to our world is OpenChain, which is an open source compliance standard. It doesn't prescribe how to run a compliance program. It says a compliance program should have certain elements — tell us what yours are. You self-certify. If someone asks, you hand over your artifacts and they decide if they're comfortable.

These management standards matter more than you'd expect, because governments reference them. A lot of times a government procurement policy will say "you must comply with ISO 1234" for data center operations, and what people don't always realize is that we can help write ISO 1234 to make sure it reflects our practices.

Management system standards also do real work as soft law — the layer of obligation that operates alongside formal regulation rather than through it. Three functions stand out. First, regulatory compliance: a regulator can require "implement an ISO/IEC 27001 information security management system" instead of writing the controls itself, which is faster, more durable, and easier to update than primary regulation. Second, efficient contracting: when both sides of a deal point to the same management standard, the diligence, the audit rights, and the remediation obligations all collapse into a known quantity instead of a bespoke negotiation. Third, trust at the scale of the economy: when millions of transactions depend on counterparties handling data, code, or supply chains responsibly, no individual contract can verify that — but a recognized management standard, with third-party certification behind it, can. That's how trust gets manufactured at scale, and it's why these standards end up mattering far more than their unglamorous reputation suggests.

Interoperability standards are where I've spent most of my career. These describe how two pieces of software interact — the format, the protocol, the interface — without prescribing how either side does the actual work.

Here's a simple example. Say you have a number-sorting standard. One side takes a series of numbers, separates them by commas, ends with a semicolon, and passes them to the other side. The receiving side sorts them lowest to highest, formats the result, and sends it back. The standard describes the communication format. It doesn't describe the sorting algorithm. You might optimize for speed. Someone else optimizes for accuracy with prime numbers. Both implementations comply with the standard as long as the inputs and outputs match.

This distinction matters enormously for patent purposes, and we'll get into that. But the core concept is that a standard describes the interface, not the implementation. There are exceptions — encryption algorithms are one, where you can't describe the standard in English alone and you need the actual algorithm — but those are outliers.

One framing point worth flagging early, because it shapes a lot of downstream confusion: an interoperability standard sets the floor for innovation and product differentiation, not the ceiling. The standard tells you the minimum you have to do to interoperate. It says nothing about what you're allowed to build on top.

This gets misunderstood often, especially in procurement and regulation. There are too many examples of policies that treat the mandated standard as the universe of permitted behavior — products that conform are "legal" or "allowed," anything else is "out of policy" or "non-compliant." That framing fails for two reasons.

First, it fails operationally. Organizations don't want their roadmaps capped by a procurement clause, so they immediately negotiate exceptions. The exceptions multiply, and within a few years the policy is more carve-out than rule. The standard ends up doing none of the work it was supposed to do.

Second, it fails economically. No economy has succeeded by using standards to set the maximum allowed behavior. The Soviet Union tried it. China tried it through the 1990s and into the 2000s. In both cases, standards-as-ceiling produced exactly the stagnation you'd expect. China's current "socialist market economy" approach is instructive: standards are still central, but they're used to set technical regulation based on flexible behavioral requirements rather than fixed implementations. The state still gets to define what "good" looks like, but innovation above the floor is allowed — encouraged, even — because the alternative doesn't work.

The lesson generalizes. When you see a procurement policy, a regulation, or a contract clause that treats an interoperability standard as a ceiling, that's a drafting error worth flagging. The standard is the common language. What gets said in it is up to the speaker.

1.2 Why Standards Matter

Standards create markets. That's the simplest version of it.

Consider the consumer electronics industry in the early 2000s. When DVD players first came out, they could cost four, five, six, seven hundred dollars. Once the market matured and prices dropped, the entire upgrade cycle was done. Consumers had their players, their TVs were fine, and nobody was buying new equipment. The industry needed the next generation to juice the cycle again. That next generation — Blu-ray — was built on standards. Without agreed-upon standards for the disc format, the compression codec, the content protection, and the display interface, you don't get an ecosystem. You get fragmented, incompatible products that consumers won't buy.

Standards also solve coordination problems. The Japanese power grid is a surprisingly good example. In the 1890s, Tokyo's electric company imported 50Hz generators from AEG in Germany, while Osaka's imported 60Hz generators from General Electric in the United States. Nobody coordinated, and both regions built out their grids independently. The voltage is the same across Japan, but the frequency difference means the two grids can't directly interconnect. Your devices can adapt — most modern electronics handle both — but the infrastructure itself cannot. When the Fukushima earthquake hit in 2011, Japan couldn't route power from the unaffected western grid to the eastern grid because the two systems were fundamentally incompatible. One standard, adopted a century earlier, would have prevented that.

The natural follow-up question is who should do the coordinating, and how. The European experience with two roughly contemporaneous mandates is instructive, because the Commission ran the same playbook twice in the same era and got opposite results.

GSM is the success story everyone cites. In the late 1980s, the European Commission — working first through CEPT and then through the newly formed ETSI — pushed a single digital cellular standard onto the European market. The mandate worked. By the mid-1990s GSM was the dominant European standard, by 2000 it was the dominant global standard, and the European cellular industry — Nokia, Ericsson, the carriers — was for a decade or so the most important player in mobile. The centralized model picked the right horse, locked it in fast, and built an industry around it.

HD-MAC is the failure story almost no one remembers. In the same period, with the same instinct, the Commission mandated a European HDTV standard for direct-broadcast satellite — analog, hybrid, designed by Philips, Thomson, Bosch, and Nokia under the Eureka 95 project, and backed by two directives (86/529/EEC and 92/38/EEC) and an Action Plan worth several billion ECU. Two things went wrong. First, the market routed around the mandate: Sky launched on the medium-power Astra satellites, which weren't covered by the directive, broadcast in PAL to the TVs consumers already owned, and absorbed the licensed competitor (BSB) within two years. Second, the technology bet was wrong: by 1993 it was clear that digital HDTV (the work that became ATSC in the US) was going to overtake analog HDTV entirely. The Commission abandoned the mandate that year and pivoted to DVB — voluntary, industry-led, digital, and modular — which went on to become the dominant digital broadcasting family worldwide.

The honest lesson isn't that mandates work or that they don't. It's that mandates work when the technology call happens to be right, and they fail catastrophically when it isn't — and the same property that makes them powerful (lock-in, no exit) is what makes the failure mode so costly. The decentralized model is messier in the short run and almost always slower to coordinate. It also doesn't have the catastrophic-failure mode, because no one is forced onto the wrong horse. Europe got HDTV right the second time around by going decentralized and digital. The pattern is worth remembering whenever someone cites GSM as proof that mandates work.

And standards carry legal and trade significance. Under WTO rules, international standards receive preferential treatment over national ones. The idea is that if we want telephone systems to work worldwide, everyone should adopt international telephone standards rather than national ones. Countries are strongly encouraged to use international standards unless there's a good reason not to. This gives international standards real power — they aren't just technical documents. They're trade instruments.

That said, the technology sector is a bit different from aerospace or manufacturing, where nearly every screw is an international standard. We move too fast. Government procurement policies don't constrain us the way they constrain the aircraft industry. Something like HTTP isn't an international standard in the formal sense, but no government is going to refuse to use it. The distinction between formal international standards and widely-adopted industry standards matters less in our world than in others. But it's still worth understanding, because when it does matter, it matters a lot.

Market-making, coordination, and trade are the headline reasons standards matter. They aren't the only ones. Several other functions show up often enough that it's worth naming them.

Trust that enables efficient contracting and diffusion. A recognized standard, especially one with third-party certification behind it, is a substitute for bespoke diligence. Two parties don't have to negotiate from scratch about what good security looks like, what good open source compliance looks like, or what good environmental practice looks like — they can point at the standard and move on. That cuts transaction costs and, just as importantly, accelerates diffusion. People adopt a technology faster when they don't have to verify it themselves.

Market access — both barrier and enabler. A standard that aligns with how a market already operates lowers the cost of entry. A standard that doesn't can become a barrier, sometimes deliberately. Regional standards are sometimes designed to favor incumbents or local industry, and the line between "legitimate technical requirement" and "non-tariff trade barrier" is blurrier than people assume. The flip side is that adopting a widely-recognized international standard is often the fastest way for a new entrant to credentialize into a market it doesn't yet have relationships in.

Consumer protection. Some standards exist to set the floor on what consumers are entitled to assume — about safety, accessibility, data handling, privacy, or product behavior. The point isn't to prescribe how a product is built. The point is to ensure that when something is sold as compliant, the consumer can rely on what that means.

Technical criteria to demonstrate conformity to new laws. As regulation moves into areas like AI, cybersecurity, and data protection, regulators increasingly need technical specificity that statutes and rules can't easily provide. Standards fill that gap. The EU AI Act, for example, leans heavily on harmonized standards to translate broad legal requirements into testable technical criteria. The standard becomes the operational definition of compliance — which makes the people who write the standard quietly very influential.

These functions overlap with each other, and with the soft-law role of management standards discussed earlier. The point isn't to draw clean lines. It's to recognize that "standards create markets" is the headline, not the whole story.

1.3 The Organizational Landscape

If you use the wrong terminology in a room full of standards professionals, you lose credibility immediately. A quick map.

Big-S Standards come from international organizations: ISO, IEC, and ITU (specifically ITU-T and ITU-R). Standards professionals will tell you these are the only "real" standards. Everything else is a small-s standard — technically a standard, but not a formal international one.

The distinction is not purely semantic. A Big-S international standard carries treaty-level obligations, government procurement preferences, and regulatory weight. A small-s standard from W3C or OASIS doesn't have that by default — but it can. A consortium spec can run through the PAS (Publicly Available Specification) process and come out the other side with an ISO stamp.

De jure standards are those that some official body has blessed. A **de facto** standard is something like the Win32 API — it became a standard by market dominance, not by process.

Below the international level you'll encounter national standards bodies (ANSI in the U.S.), regional European bodies (CEN, CENELEC, ETSI — backed by considerably more government involvement than anything you see in the U.S.), and a large ecosystem of industry consortia and foundations where most technology standardization actually happens. Chapter 2 walks through each category in detail. The point for now is just that the category of organization shapes everything downstream: governance rules, IPR defaults, confidentiality expectations, and the route to government adoption.

1.4 Where the Work Gets Done

Day to day, most of the standards work that matters to technology practitioners happens in industry consortia and foundations — IETF, W3C, OASIS, the Linux Foundation, the Joint Development Foundation, and dozens of smaller bodies. These are the workhorses, and Chapter 2 covers them in depth.

Here's an analogy that's worth carrying through the rest of the book. Think of a standards organization as a warehouse.

The walls, floor, and roof are the legal framework — the bylaws, the IPR policy, the antitrust guidelines, the membership terms. They define the building. People bring buckets of ideas in through the door; those are contributions. Inside, there are tables where experts gather and the buckets get emptied out and the contents are debated. Those tables are the working groups. Some people stand in the room without bringing buckets — they listen, they react, they shape the conversation without contributing material themselves. Those are participants. When the work is done, the result gets stacked on the loading dock at the back of the warehouse. Implementers walk up, take a copy, and go build something with it.

The roles aren't exclusive, and forgetting that is a common source of confusion. A contributor is usually also a participant, and is often an implementer of the spec they helped write. A participant who never contributed material can still go pick something up off the loading dock and build with it. And the people who walk up to the loading dock — the implementers — generally aren't contributors or participants at all. Most implementers of any successful standard never set foot inside the warehouse. That's the whole point of having a loading dock.

Four roles, three groups of rights. Contributors have rights in what they put into the bucket. Participants have rights to be in the room and to be heard. Implementers have rights in what they pick up off the dock. The whole job of a standards organization — the bylaws, the IPR policy, the governance — is to keep those three sets of rights in balance. Tilt too far toward contributors and you get a standard nobody can implement. Tilt too far toward implementers and nobody contributes. Tilt too far toward participants and the work doesn't ship. Most of the choices we'll work through in the rest of this book — patent policy, scope, exclusions, governance, voting, copyright — are choices about where on that triangle a particular organization wants to sit.

One more thing worth flagging up front, because it shapes how you should think about every battle inside the warehouse: just because you win the standards vote doesn't mean people will implement the standard. Most interoperability standards are voluntary. There is no regulator forcing anyone to walk up to the loading dock. The implementers come — or don't — based on whether the spec actually solves a problem they have, whether they trust the process that produced it, and whether the people whose code they need to interoperate with are also going to ship it. If you push a standard through to publication by out-voting the people whose buy-in you actually needed, you can win the battle and lose the war. The loading dock stays empty. This is why so much of the practitioner's job is consensus-building during development rather than vote-counting at the end. The vote is the receipt; the consensus is the product.

One durable pattern to internalize before we move on: as soon as a standards organization becomes an incorporated entity with staff, it develops its own agenda. That may or may not be a problem, but it is always a factor. If all you need is a lightweight home for some specs, you want to understand whether the organization's institutional interests align with yours. Organizations behave like organizations, not like neutral hosts.

1.5 Standards and Open Source: The Convergence

This is the trend that changed how we think about all of this.

For most of its history, the standards world and the open source world operated in parallel and didn't interact much. Standards had their formal processes, their patent policies, their governance structures. Open source had licenses, contributor agreements, and community-driven development. The IETF's principle of "rough consensus and running code" bridged the two cultures early — IETF standards development often relied on open source implementations as proof of concept — but outside of IETF, the two worlds were largely separate.

Then they started doing the same thing.

The HTML5 story is the clearest example. HTML4 was a W3C standard. When it came time for the next version, W3C's staff believed the future was the "semantic web" — markup that described not just layout but meaning, making it easier for machines to understand web content. Academically interesting. But a group of browser vendors — Apple, Mozilla, and Opera — just wanted a better markup language. In 2004, they formed the WHATWG (Web Hypertext Application Technology Working Group), forked HTML, and started iterating on it as a living standard. Google and Microsoft joined later. Meanwhile, W3C's own effort — XHTML 2.0 — went nowhere.

The result is a strange arrangement that persists today: the actual HTML work happens at WHATWG, but it gets published through

W3C so W3C still looks like the steward. The crown jewel walked out the door, and the institution had to adapt.

This matters because it demonstrates something fundamental: if people don't like what a standards body is doing, they will fork the work. It's harder than in open source, and the ecosystem isn't designed for it, but it happens. And increasingly, the fork destination is an open source project rather than another standards body.

At Microsoft, we saw this convergence coming. We had always had a separate standards team and open source team. The two groups were sizable and operated independently, with different cultures and different risk profiles.

Eventually, we combined them. The logic was straightforward: open source was increasingly being used in place of standards. It's a different kind of plug fitting into the same outlet. If we wanted to advise our product teams on how to do interoperability — which is ultimately what both standards and open source solve — we needed to be able to look at the full picture and recommend the right tool for the situation. Sometimes that's a formal standard. Sometimes it's an open source project. Sometimes it's both.

That convergence is now the norm across the industry. The Joint Development Foundation supports both standards and open source governance. Organizations like the Linux Foundation host both traditional specifications and open source projects. The line between "standard" and "widely-adopted open source project" has blurred to the point where the distinction is often more about process and IP terms than about the output itself.

For attorneys advising on these engagements, this means you need to understand both worlds. A standards patent policy and an open source contributor license agreement are solving related but different problems — and they cover different things. An open source license covers the specific code (or text) that contributors submit. It does not cover the specification as a whole. A standards patent commitment covers necessary claims across the entire specification regardless of who contributed what. That distinction becomes critical when someone releases a "spec" under an open source license and treats that as patent coverage; it isn't. Chapter 5 develops this point in depth. For now, the foundation to build from is this: standards and open source are two paths to the same destination — interoperability — but they reach it through different legal machinery. Know both.

1.6 Key Terminology Quick Reference

Before we move on, here's the shorthand you'll encounter throughout this book.

Term	Meaning
SDO	Standards Development Organization — broad term covering most standards bodies
SSO	Standards Setting Organization — sometimes used for the "big" international bodies specifically
Consortium	Usually a smaller, often contractually based collaboration; may or may not be incorporated
Foundation	Generally an incorporated entity (often 501(c)(6)) that hosts standards or open source work
SIG	Special Interest Group — a small, focused collaboration; the term has fallen out of favor
Big-S Standard	A standard from an international body (ISO, IEC, ITU)
Small-s standard	An industry standard from a consortium or foundation (W3C, OASIS, IETF, etc.)
De jure	A standard blessed by an official body through a formal process
De facto	A standard by market adoption, not formal process (e.g., Win32 API)
PAS	Publicly Available Specification — a process for submitting consortium specs to ISO/IEC for international standardization
ANSI	American National Standards Institute — accredits U.S. standards bodies

IPR Policy	Intellectual Property Rights Policy — the patent and copyright rules governing a standards body; the subject of most of this book

Use the right terms. Standards professionals notice when you don't.

Practice Tips

1. Know which type of standard you're dealing with — health/safety, management/process, or interoperability — before advising on IP strategy.
2. Learn the difference between Big-S and small-s standards — formal international standards carry treaty-level weight that consortium standards don't.
3. Evaluate both standards and open source as paths to the same interoperability goal.
4. Map the relevant standards bodies for your sector early — the organization determines the governance rules.
5. Use correct terminology — standards professionals notice immediately when you don't.

Chapter 2 — The Organizational Landscape: Types of Standards Bodies

The first question in any standards engagement isn't about IP policy or patent commitments. It's about structure. Where is this work going to happen? In an existing organization or a new one? Under a corporate entity or a multi-party contract? With staff and dues, or on a volunteer basis?

These structural choices cascade into everything else — governance, IP terms, membership dynamics, decision-making power. Get the structure right and the rest follows. Get it wrong and you'll be working around the problems for the life of the project.

This chapter walks through the organizational landscape from the top down, then offers a practical decision framework for choosing a home for new work.

2.1 Formal International Organizations

At the top of the hierarchy sit three international standards bodies: ISO, IEC, and ITU (specifically ITU-T and ITU-R).

ISO and IEC collaborate on technology standards through a joint structure called JTC-1 (Joint Technical Committee 1), which is where work like MPEG video codecs, JPEG image formats, and other technical specifications gets done. Below JTC-1, the work is further divided into subcommittees — SC29 for media coding, for instance. Each level has its own procedures, its own delegates, and its own politics.

ITU-T is a UN-chartered organization focused on telecommunications. It's the most formal of the three, with strong ties to national regulatory bodies.

One critical feature of the international system: companies don't participate directly. You participate through your **national body**. In the United States, that's INCITS for JTC-1 work, operating under ANSI accreditation. National bodies develop a national position, and that position is cast as a single vote. The US vote counts the same as any other country's vote, regardless of the size of the delegation behind it.

This creates interesting dynamics. The same company may be deeply involved in the technical work, even chairing a committee, but the vote that matters is the national body's. During contentious ballots, companies have to lobby their own national body — and sometimes carry votes they disagree with. It's diplomacy, not direct democracy.

ISO, IEC, and ITU-T all share a common patent policy framework, though each has different implementation guidelines. Think of it as one text with different interpretive traditions. The policies themselves seldom change, and for the most part they're workable. The differences tend to be procedural — things like what information you must provide when declaring a patent, or how the disclosure process operates.

Standards from these bodies carry the most formal weight. WTO agreements give international standards preferential treatment over national ones in trade disputes and government procurement. That matters more in some industries (aerospace, telecom) than in ours, but when it matters, it matters a lot.

The PAS Transposition Process

One mechanism that bridges the international and consortium worlds is the **PAS (Publicly Available Specification) submission process**. This allows an organization that has been granted PAS submitter status — W3C and the Joint Development Foundation both have it — to develop a specification through its own process and then submit it to ISO/IEC JTC-1 for an international vote. If the national bodies approve it, the consortium spec becomes a formal international standard.

This is significant because it lets you develop a spec where you want — in a faster, more technically focused environment — and then layer the international formality on afterward. The PAS process avoids the overhead of developing the spec within the international system while still achieving the benefits of international recognition: treaty-level status, government procurement eligibility, and regulatory weight.

2.2 National and Regional Standards Bodies

Below the international level, national standards bodies develop standards for their home markets and contribute to international processes.

ANSI (American National Standards Institute) does not write standards. It accredits other organizations to develop them. An ANSI-accredited standard carries the designation "American National Standard," which is valuable for US government procurement. ANSI also coordinates the US participation in ISO and IEC.

In **Europe**, three organizations fill a similar role: CEN, CENELEC, and ETSI. The European model is more directive than the American one. Where the US tends to let industry self-organize, Europe is more likely to say: we need a standard for this technology, ETSI go develop it, and we'll have regulatory input into the process.

ETSI is particularly significant because it develops the mobile telecommunications standards as part of the broader 3GPP framework. If you're working in anything touching cellular networks, ETSI's rules and FRAND policies are inescapable.

A note on how that "regulatory input" actually plays out lately. The traditional picture is that the Commission issues a standardisation request, the European Standardisation Organisations develop the standard, and the Commission stays out of the room while the work is done. That picture is getting fuzzier, particularly around the development of harmonised

European standards (HENs) — the standards that, once cited in the Official Journal, give implementers a presumption of conformity with EU law.

In recent work, Commission representatives have been physically in the room. They have been stating "red lines" — positions the standard cannot cross if it is to be cited as a HEN — but declining to be quoted in the minutes. The substantive influence is real; the documentary trail is thin. For participants and counsel, that creates a practical problem: you are negotiating against constraints that aren't in the policy documents, can't be cited back, and may not be visible to anyone who wasn't in that meeting. It also raises questions about whether the standardisation process retains the independence and transparency that the harmonised-standards regime is supposed to depend on. This is a live tension, and it's worth watching as the case law and the Commission's own practice continue to evolve.

In **China**, the standards system is more state-directed than either the US or European models, and it has become impossible to ignore in any serious global standards strategy. The realistic picture of global standardisation today is a triangle — the United States, the European Union, and China — and any framing that leaves China out is incomplete. **SAC** (Standardization Administration of China) sits at the top as the national standards body and the formal interface to ISO and IEC. The two organisations most counsel and standards professionals actually deal with at the working level are **CESI** (China Electronics Standardization Institute), which covers ICT and electronics standardisation, and **CATR** (China Academy of Information and

Communications Technology), which covers telecommunications. Both are technically research institutes sitting under the relevant ministries, but in practice they drive much of the substantive work and shape the national positions that get carried into international forums.

The Chinese system also uses mandatory national standards — the GB series — as a market-access lever in a way the American and European systems generally do not. The GB18030 character-set example in Chapter 3 is the canonical illustration. The strategic implication for industry is one we'll come back to in Chapter 13: the international route through ISO/IEC is frequently the practical path to a Chinese national-standard adoption, because consortium specifications on their own have limited standing inside the Chinese system. Counsel advising on a global standards strategy in any consequential technology area should assume China is part of the picture from the start, not a follow-on consideration.

Accreditation— whether ANSI accreditation in the US or recognition by European authorities — comes with obligations. Accredited organizations must follow due process requirements: openness, balance of interests, consensus-based decision making, the right to appeal, and transparency. These requirements exist to ensure that standards reflect broad input rather than the interests of a few dominant players. They also provide antitrust protection — a standard developed through a recognized due process is far less likely to face antitrust challenge than one developed behind closed doors. We'll cover due process in depth in Chapter 4.

The differences across the American, European, and Chinese approaches have practical consequences for how you engage, who you lobby, and what rules apply. It's worth understanding all three even if most of your work falls inside one of them.

2.3 Industry Consortia and Foundations

This is where most day-to-day technology standards work happens.

IETF (Internet Engineering Task Force) develops the core Internet protocols — TCP/IP, HTTP, TLS, DNS. It operates under the Internet Society and has a distinct culture: no formal membership, no voting in the traditional sense, and a principle of "rough consensus and running code." IETF participation is individual, not corporate, which creates a different dynamic than most standards bodies.

W3C (World Wide Web Consortium) develops web standards — HTML, CSS, Web APIs, and related technologies. It incorporated as its own legal entity in 2023 after years of operating under a host institution model. W3C's royalty-free patent policy, first adopted in 2004 and updated several times since, was one of the first in the industry and remains influential. Participation is organizational, with membership fees that vary by organization size.

OASIS (Organization for the Advancement of Structured Information Standards) focuses on enterprise, security, and interoperability standards. It offers multiple IPR policy modes —

RAND, royalty-free on RAND terms, and royalty-free on limited terms — giving technical committees flexibility to match IP terms to their context. OASIS has a well-defined staged specification process and is particularly strong in areas like cybersecurity and cloud standards.

The **Linux Foundation** operates differently from these. It's primarily a hosting organization — it provides infrastructure, operational support, and governance frameworks for a wide portfolio of open source projects and, increasingly, standards initiatives. The Linux Foundation doesn't develop standards itself; it hosts projects that do. Its scale is significant: hundreds of projects, substantial revenue, and a large professional staff. For projects hosted under the Linux Foundation, it provides accounting, legal support, event management, and program management. The tradeoff is that hosted projects operate within the Linux Foundation's rules and fee structures.

These organizations have open membership, technology-driven cultures, and established processes. If you're doing web work, W3C is the natural home. Internet protocols go to IETF. And so on. Each has its own governance model, its own IP policy, and its own institutional personality.

One thing to watch for: people develop strong affinities for the organizations they know. A standards professional who has spent years working in W3C will naturally want to bring new work there, because they know the people, they know the rules, they know how to navigate the hallways. That's not always wrong — institutional knowledge has real value — but it means you should evaluate the recommendation independently. Is this the right venue for the technology, or is it just the venue the proposer is comfortable with? As counsel, be aware that the recommendation to "go to organization X" may be serving the proposer's professional interests as much as the company's.

Working Within an Existing Organization

Joining an existing standards body is the path of least resistance. You don't need to set up a corporate entity, file for nonprofit status, establish bank accounts, or hire staff. You're joining a gym — you use their equipment, hold meetings in their facilities, and follow their rules. The tradeoff is that you're subject to those rules, including IP policies and governance structures you didn't write and may not love.

This approach makes the most sense when the work fits naturally into an existing organization's scope, when you don't need your own identity or brand, and when you're comfortable with the organization's process. A lot of standards work fits these criteria perfectly. Not everything needs a new foundation.

Setting Up a New Organization

Sometimes you do need your own identity. You want your own name, your own website, your own brand. You want to control membership criteria. You need a certification program or a logo. Or you want to work with a specific set of partners and don't want the fully open participation model that established organizations require.

The decision to form a new organization is significant, and the single most important question in making it is deceptively simple: **do you need to hold money?**

If the answer is no — if participants can volunteer their time, rotate meeting hosting, and the work doesn't require paid staff — you can stay lightweight. A contractual consortium or the Community Specification License (discussed in Chapter 9) may be sufficient.

If the answer is yes — you need to hire a program manager, pay for a website, organize conferences, run a certification program — you need a corporate entity with a bank account. And that changes everything.

2.4 Contractual Consortia

A contractual consortium is the simplest multi-party structure. There's no incorporated entity. It's a contract — everyone signs the same agreement, and the rights and obligations flow from that agreement alone.

These have real advantages. They're cheaper to set up and cheaper to run than incorporated entities. There's no corporate formality, no annual filings, no nonprofit status to maintain. For a small group working on a focused problem, they can work well.

The disadvantages are equally real.

They don't scale. A contractual consortium works for maybe eight to twelve companies. Beyond that, the logistics of getting everyone to sign and managing the agreement become unwieldy.

They're hard to amend. If you want to change the agreement, you generally need everyone to re-execute. That effectively gives every participant a veto. In practice, changes can take months, particularly when participants with complex internal approval processes are involved.

They can't hold money. There's no entity to maintain a bank account. If you need to host a meeting, someone volunteers and pays. Need a website? Someone builds one on their own time. Need a program manager? You don't have one. This volunteer model works until it doesn't, and when it stops working, there's no easy fix.

They have longevity problems. Contractual consortia tend not to die formally. People lose interest and walk away, but the agreement just sits there. Meanwhile, the spec is still being implemented. Years later, someone needs an authoritative copy and there's no organization maintaining one. The spec ends up floating on a random website with no clear provenance.

Contractual consortia can transition to formal incorporated entities, and some have done so successfully. But the transition itself is a heavy lift — it requires renegotiating the underlying agreements, establishing corporate governance, and managing the administrative overhead of incorporation, all while trying to keep the technical work moving forward. The distraction can be significant. The evolution of frameworks like the Joint Development Foundation has provided an increasingly viable and preferred alternative, offering the benefits of a formal structure without the startup burden of building one from scratch.

2.5 Incorporated Standards Organizations

When you need staff, money, and formal governance, you incorporate. That means choosing a state of incorporation, filing for nonprofit status (typically 501(c)(6) for industry standards), establishing bylaws, setting up bank accounts, and managing ongoing corporate formality.

Most organizations outsource the operational side to **association management companies** — firms that specialize in running standards bodies. They provide accounting, program management, website maintenance, and event logistics. The advantage is that you don't have to build that capability yourself. The disadvantage is that these companies are businesses with their own revenue incentives. They'll push for additional services, strategic consulting, expanded programs. You need to manage them actively.

The alternative is to operate under a larger umbrella organization. Several options exist: the Linux Foundation, IEEE ISTO, the Joint Development Foundation, and OASIS all provide hosting infrastructure for standards projects. Each has its own fee structure, operational model, and degree of autonomy for hosted projects. Choosing among them involves evaluating cost, flexibility, institutional fit, and the community you want to attract.

Membership structures in incorporated organizations typically involve tiered classes — a steering committee or board with greater decision-making power, and one or more general membership tiers with participation rights but less control. The design of these tiers matters enormously. The top tier typically controls the organization's direction, approves specifications, and manages the budget. Lower tiers contribute technical work and participate in committees but may not have a vote on organizational governance.

Getting the membership criteria right is one of the hardest practical problems in setting up a new organization. The criteria need to be objective — antitrust law requires that similarly situated parties be treated alike. But every "objective" criterion is, at some level, subjective. If you offer a board seat to one major company in a sector, you need a principled reason for saying no to its direct competitor. Geographic diversity, industry segment representation, and revenue thresholds are common approaches, and getting the balance right requires careful thought about the specific context.

2.6 The Joint Development Foundation

The Joint Development Foundation was created to solve the problems described above — specifically, the three-to-nine-month startup period consumed by negotiating agreements that all look the same at the end.

JDF provides a **standards organization in a box**: predefined governance documents, predefined IP policy modes, and a corporate structure that's ready to use. A new project can be launched in days rather than months.

The Legal Structure

JDF uses a **series LLC** model. The parent entity is a 501(c)(6) nonprofit organized in Washington State. Underneath it sits JDF Projects LLC, a Delaware Series LLC. Each project operates as a separate series within JDF Projects LLC. This provides structural insulation — the liabilities and IP commitments of one project don't bleed into another. A company participating in Project A has no patent exposure to Project B.

This structure was novel when JDF introduced it. It was borrowed from real estate law, where series LLCs are used to separate properties within a single entity. There was initial skepticism about whether it would work for standards, but it has proven effective. The Linux Foundation subsequently adopted the same series structure for its own projects.

How It Works

When a new project is chartered under JDF, the participants select from predefined governance terms and IP policy modes — including multiple patent licensing options. The working group charter specifies which mode applies. There's no negotiation over the framework itself. The framework is non-negotiable by design.

This is a deliberate philosophical choice. By removing the legal negotiation from the project setup, JDF keeps attorneys out of the critical path. Engineers can focus on the technical work without waiting months for lawyers to agree on terms that, historically, end up within the range of reasonable anyway.

The tradeoff is flexibility. You can't customize the JDF framework the way you can with a bespoke agreement. For most projects, that's fine — the predefined options cover the common cases. For projects with unusual requirements, JDF may not be the right fit.

Operational Reality

JDF projects operate under the Linux Foundation's infrastructure. This provides accounting, tax filing, program management, and administrative support. The advantages are significant: no startup cost for operations, access to established processes, and a recognized institutional home.

JDF has become the default starting point for new standards collaborations in many areas. It now hosts dozens of active projects. For most engagements where the predefined framework fits, the speed and simplicity of JDF outweigh the limitations.

2.7 The Decision Framework

When someone comes to you and says "we want to do a standards thing," here's the decision tree.

Do you even need an organization? Sometimes a unilateral spec release under an open license is enough. If you're just publishing documentation for others to implement, you may not need collaboration infrastructure at all.

Do you need to collaborate? If yes, but the collaboration is lightweight — a handful of parties, no money needed, no formal identity required — tools like the Open Web Foundation agreements or the Community Specification License (CSL) may be sufficient. They provide IP terms without organizational overhead. We'll cover CSL in more detail in Chapter 9.

Can the work fit in an existing organization? If the technology aligns with an established body's scope, and you don't need to control membership or branding, joining an existing organization avoids the cost and complexity of forming a new one. This is often the best answer, even if it's not the most exciting one.

Do you need your own identity? If you need a name, a brand, a website, a certification program — or if you need to control who participates and at what level — you're forming a new organization.

Do you need money? This is the fork in the road. If no, a contractual consortium or similar lightweight structure can work. If yes, you need a corporate entity with a bank account, and the overhead that comes with it.

How fast do you need to move? If speed matters — and it usually does — JDF or a similar predefined framework avoids months of legal negotiation. If you have genuinely unusual requirements that the predefined frameworks can't accommodate, a bespoke structure may be necessary, but go in with eyes open about the time and cost.

The right answer depends on the specific engagement. But the questions are always the same. And the most important one — do you need money — is the one that most often determines the path.

Practice Tips

1. Always ask first: "Do we need to hold money?" — that single question determines the organizational structure.
2. Check whether the IP commitment is organization-wide or working-group-specific.
3. For new projects, prefer predefined frameworks (JDF, CSL) over bespoke agreements.
4. When evaluating a proposed venue, ask: "Is this the right venue for the technology, or just the venue the proposer is comfortable with?"
5. Map your client's corporate structure before joining — ambiguity around "controlled" entities creates unexpected patent exposure.

Chapter 3 — The Intellectual Property Landscape of Standards

Four categories of intellectual property come up in standards work: copyright, patents, trademarks, and trade secrets. Of these, patents get the most attention — and deserve it, because that's where most of the risk and complexity lives. But copyright and trademark each carry their own traps, and trade secrets present a unique set of concerns in a context built around sharing. Understanding how all four interact is essential before you start reading actual policies.

This chapter is the overview. Subsequent chapters will take you inside specific patent policies, necessary claims definitions, exclusion mechanisms, and the rest of the machinery. Here, the goal is to establish the framework so the details in subsequent chapters have context.

3.1 Copyright: Broader Than You Think

When someone makes a contribution to a standards body — a technical proposal, a draft section, a diagram — that contribution carries copyright. The standards body needs rights to that material. How it gets those rights varies, but the mechanism is almost always one of two things: a broad license or an assignment with a license back.

Most people prefer the license approach. It's less intrusive. You keep ownership of your contribution, which means you can submit the same material to multiple standards bodies if you're working on related specs. An assignment feels grabby, even when it comes with a broad license back. There's additional paperwork. People tend to avoid it when they can.

In practice, the distinction matters less than you'd expect. Whether it's a license or an assignment with a grant back, the rights the standards body receives are exceedingly broad. They need to be. The organization has to be able to evolve the material, create derivative works, publish it, and maintain it indefinitely. Don't assume you're retaining meaningful control over contributed material. The license is that broad by design.

Inbound and Outbound Copyright

It's useful to think about copyright in standards as having two directions.

Inbound copyright is the rights the standards body receives from contributors. This is the license or assignment discussed above — the mechanism that allows the organization to assemble, edit, publish, and maintain the specification.

Outbound copyright is the rights the standards body grants to implementers and the public. This is how the finished specification gets distributed. Some organizations sell copies, some publish under Creative Commons licenses, some make specs freely available on their websites. The outbound terms determine who can access the spec, whether they can reproduce it, and whether they can create derivative works.

The inbound and outbound sides can be asymmetric. A standards body might receive broad inbound rights from contributors but distribute the finished spec under restrictive outbound terms — for example, allowing implementers to read and use the spec but not reproduce or modify it. Others may be broad in both directions. The combination of inbound and outbound terms shapes the practical accessibility of the standard.

Who Owns the Spec?

Here's where it gets interesting. In most cases, the standards body doesn't own the entire specification outright. What it owns — or more precisely, what it has rights to — is the collective work. Individual contributions remain with their contributors. The standards body assembles them into something that functions as a unified document, and the rights in that compilation are what gives the organization its copyright position.

This works well enough when you have an incorporated standards body sitting in the middle. But in a contractual consortium — one with no corporate entity — there's no one to hold those collective rights. What you end up with is effectively a massive cross-license among all the participants. Everyone grants rights to everyone else. It works, but it's a different legal posture than having a central entity that can enforce the copyright in the spec.

Why does enforcement matter? Because for many traditional standards bodies, selling the spec is the business model. ISO, IEC, and similar organizations charge for access to their standards — sometimes hundreds of dollars per document. They need the ability to prevent unauthorized reproduction and distribution. The copyright in the collective work is what gives them that.

The Access-to-Law Problem

This business model is under pressure, and it's worth understanding why.

When a government normatively references a standard in legislation or regulation — building codes are the classic example — citizens are effectively required to comply with a document they may have to pay to access. You want to build a house in Redmond? You need to meet the building code. But the building code incorporates standards developed by private organizations, and those organizations charge for copies.

The democratic objection is straightforward: why should citizens have to pay a private entity to access the law? Some standards bodies have responded by allowing limited public viewing — you can read the spec in a reading room, but you can't take a copy home. Recent court decisions, particularly in Europe, have pushed further toward requiring free access when a standard is incorporated into law.

For most technology companies, this is not their fight. It's a business model problem for standards bodies. The preference is for all specs to be publicly available, but that doesn't mean it's wise to pick a fight with a standards body over their revenue model.

Where this has gotten more acute recently is around AI training. Some standards bodies have started posting notices on their websites explicitly prohibiting the use of their standards as AI training data — ISO updated its Terms of Use in 2023 to address this directly. The practice is emerging, not universal, but at least one organization has sent demand letters to technology companies. Their concern is that if someone can ask an AI to explain how to comply with a standard, there's no reason to buy the document. That concern has some basis, even if the practical risk of relying on AI-generated building codes seems obvious.

Creative Commons and Forking Risk

Many technology-focused standards bodies publish their specs under Creative Commons licenses or even on open wikis. The philosophy is that broad access leads to broad adoption.

The criticism — and I think it's valid — is that this creates forking risk. If someone can modify a spec on a wiki, and you build a product to that modified version, interoperability breaks. In practice, this hasn't been a significant problem. But it's a good reminder to advise your clients to always pull specs from the canonical source.

Copyright as a Stability Mechanism

It's worth stepping back and recognizing that copyright in standards isn't just about revenue or enforcement. It serves a structural function: maintaining the integrity of the spec over time.

A canonical, copyright-protected specification gives implementers confidence that the document they're building to is the real one. It prevents unauthorized forks that would fragment the ecosystem. When someone modifies a spec — even with good intentions — and products get built to the modified version, interoperability breaks. Copyright is the legal mechanism that keeps a single authoritative version in place.

This is one of the reasons that even organizations philosophically committed to openness still maintain copyright control over their specs. You can make a spec freely available and still prevent unauthorized modifications. The Creative Commons Attribution-NoDerivatives license, for instance, allows free distribution while prohibiting forks. The goal is broad access with long-term stability — and copyright is the tool that makes that possible.

3.2 Patents: The Main Event

Patents are where the complexity, the risk, and most of the case law lives. The entire second part of this book is devoted to patent policies, but here's the conceptual foundation.

The Monopoly Problem

Start with first principles. If you have a patent that reads on a standard, and that standard is widely adopted, you have a monopoly position.

In a normal market, patents create leverage but not unchecked power. If someone holds a patent on a particular technology, you can design around it, license a competing patent, or build something different. Market forces impose some discipline on what a patent holder can charge.

Standards eliminate that discipline. When a technology is embedded in a standard, every implementer must use it. There is no designing around it. The standard says: do it this way. If you hold a patent on "this way," you can charge whatever you want. You can seek injunctions. You have monopoly rents on a captive market.

This is the problem that every patent policy in every standards body is trying to solve.

RAND and FRAND

The most common solution is the RAND commitment — Reasonable and Non-Discriminatory licensing terms. In Europe, it's called FRAND, with the F standing for Fair. They mean the same thing, and no, the F does not mean Free. This is a common and consequential misunderstanding.

A RAND commitment is a commitment to license — not an actual license. The distinction matters. When a patent holder makes a RAND commitment, they are promising to offer a license on reasonable and non-discriminatory terms to anyone who wants to implement the standard. But the actual license itself — the specific terms, the rate, the scope — is left to the patent holder and the implementer to negotiate bilaterally, outside the standards body. The standards organization facilitates the commitment. It does not broker or administer the resulting licenses.

The idea is to restore something approximating competitive market conditions. The patent holder can still charge for necessary claims, but they can't leverage the monopoly created by the standard to extract monopoly pricing.

What constitutes "reasonable" is, predictably, where all the litigation happens. The policy terms themselves are rarely contested. It's the rates that get fought over. And those fights are substantial — they involve hundreds of millions of dollars in the wireless and video codec spaces. There's also an ongoing debate about whether a RAND commitment prohibits the patent holder from seeking injunctions against implementers. Opinions vary by jurisdiction and by policy, and courts have not been uniform on this point. It's one of the more contested open questions in standards patent law.

Necessary Claims

The RAND commitment doesn't apply to your entire patent portfolio. It applies only to what are called **necessary claims** (sometimes "essential claims") — patent claims that you cannot avoid when implementing the standard.

This concept is the heart of every patent policy and deserves its own chapter, which it gets. But at the overview level, the logic is this: if a patent claim is truly unavoidable — if there is no way to implement the standard without practicing that claim — then the patent holder has a monopoly, and the RAND commitment kicks in.

Counterintuitively, while disputes over essentiality do occur, they are far less frequent than disputes over royalty rates. Both sides typically have incentives to agree that a claim is necessary — the game theory favors it — and litigating essentiality is expensive enough that parties often concede the point to focus the fight on the rate. The game theory behind this dynamic is explored in Chapter 6.

Royalty-Free vs. Royalty-Bearing

Not all patent policies allow royalties. Royalty-free (RF) policies — sometimes called RAND-RF or RAND-Z (zero royalty) — require patent holders to license their necessary claims at no charge. The implementation is free; the commitment is irrevocable (usually with a defensive termination provision).

The choice between RAND and royalty-free is one of the most significant decisions in setting up a standards body, and it shapes everything downstream. Industries with large existing patent portfolios — telecom, video codecs, wireless — tend to use RAND. Software-focused and web-focused organizations tend to use royalty-free. We'll explore this divide in depth in Chapters 7 and 8.

The Certification Shortcut

One practical consequence of standards patents is worth noting here. If your product carries a certification logo — Wi-Fi, Bluetooth, whatever — the logo serves as strong evidence that the product practices any necessary claim. It's not a formal admission of infringement in a legal sense, but it dramatically simplifies the patent holder's case.

A patent holder pointing to a certified product doesn't need to prove infringement claim by claim. The logo tells the world you implement the standard. If they hold a necessary claim, the conversation moves quickly to the rate — which is one reason patent policies in standards have so much practical significance.

3.3 Trademarks and Certification Marks

Trademark in standards will feel familiar if you've worked with open source foundations. There are similar dynamics around brand control and licensing. But standards add a certification layer that creates its own set of issues.

The Value of the Logo

For many standards, the trademark — specifically the certification mark — is the most commercially powerful piece of IP the organization holds. Not the spec. Not the patents. The logo.

Retailers won't stock a Wi-Fi router without the Wi-Fi logo. It's not about legal obligation; it's about returns. A router without the logo might not interoperate with everything, and the retailer doesn't want to deal with customers bringing it back. So the logo becomes a market access requirement.

More forcefully, trademarks can be enforced at the border. Customs officers inspecting a shipment of, say, Bluetooth headsets will check whether the product is listed in the Bluetooth organization's licensed product database. If it's not, the shipment gets seized. That's a powerful enforcement mechanism — far more immediate than patent litigation.

Self-Certification vs. Formal Testing

Conformance programs come in several forms.

Self-certification is the most common. The standards body publishes a set of requirements, you test your product internally, and you declare compliance. OpenChain's open source compliance standard works this way. You describe your program, you self-certify, and if someone asks, you hand over your artifacts.

Formal third-party testing is the heavier approach. You submit your product to a testing lab — sometimes UL, sometimes a university lab, sometimes a lab run by the standards body itself — and they verify compliance before you get the logo. Some organizations use a hybrid: formal testing for your first product, self-certification for subsequent ones.

A foreign-market example worth knowing. China uses mandatory standards compliance, verified through formal testing, as a market access requirement. The leading example is **GB18030**, the national standard for the Simplified Chinese character set. Any product sold in China that handles text — operating systems, browsers, document software, fonts, embedded systems with displays — has to support GB18030, and conformance has to be demonstrated through the official testing process. There is no self-certification path. There is no "good enough" alternative. If you want access to the Chinese market, you support GB18030 and you get certified. The standard itself is unobjectionable on the merits — encoding the world's largest script is a real interoperability problem — but the certification regime is also doing market-access work, and that's the part counsel needs to understand. Several other Chinese GB standards function the same way across other product categories. When evaluating a product launch in China, the GB compliance pathway is often as important as the IP analysis.

When Certification Overrules the Spec

This is the trap to watch for. Sometimes the conformance program tests for things that go beyond what the spec requires.

A specification might describe a feature as optional — implementers can include it or not. But the conformance program might test for that optional feature and fail you if it's missing. Your product complies with the spec. It doesn't pass certification. No logo.

This has real consequences. It effectively elevates optional spec features to mandatory ones, not through the standards process but through the certification program. The Java certification dispute between Sun and Microsoft was partly about exactly this — the certification program was testing for capabilities beyond what the spec strictly required.

For antitrust reasons, most standards bodies separate their spec development and certification functions into different organizational arms, or rely on external organizations entirely. The IEEE develops the 802.11 wireless specifications. The Wi-Fi Alliance, a separate organization, runs the certification program and grants the Wi-Fi logo. This separation is deliberate and important.

Profiles

One more piece of the trademark and certification picture: profiles.

A broadly written standard may support many configurations. Different implementers may tune the settings differently, and products tuned differently may not interoperate even though they all comply with the spec. Profiles solve this by defining specific configurations for specific use cases.

MPEG DASH is a good example. The standard describes how to handle media streaming under varying network conditions — buffering, quality switching, congestion management. But the standard had so many options that different implementations set the dials differently and couldn't interoperate. The DASH Implementers Forum created profiles on top of the standard: if you're streaming on the web, use these settings; if you're on satellite, use those. This was what ultimately allowed services like Netflix to consolidate from dozens of encoding configurations down to a manageable few.

Profiles are sometimes developed by the standards body itself, sometimes by external organizations. Either way, they sit between the spec and the certification program and are often where the practical interoperability gets defined.

3.4 Trade Secrets: The Wrong Tool for Standards

Trade secrets and standards are fundamentally incompatible. Standards are built with and among competitors, and their entire purpose is to become publicly available for broad implementation. A trade secret, by definition, derives its value from being kept confidential. These two concepts don't coexist.

If you're advising someone who wants to contribute proprietary technology to a standards body while maintaining trade secret protection, the answer is almost always: don't. It is possible to share high-level, non-essential information in a standards setting and preserve trade secret status — particularly if appropriate NDAs are in place and the disclosure is carefully controlled. But the moment the inventive concept itself is shared in a multi-party setting with competitors — often dozens of them — trade secret status is effectively gone. Even if everyone in the room has signed a confidentiality agreement, the practical ability to maintain secrecy across that many parties and organizations is negligible.

More importantly, the intent of standards is publication. The spec will eventually be made available to the world. Any technology embedded in it will be described in sufficient detail for independent implementation. You cannot contribute to a standard and expect to maintain trade secret protection over the contributed material.

This doesn't mean proprietary information never enters a standards discussion. It does. Product roadmaps get referenced. Implementation strategies get discussed. Business plans come up in the context of use cases. But the onus is on the discloser to protect their own information. Standards meetings happen among competitors. If you share something confidential in that setting, that's on you.

We'll discuss confidentiality provisions in standards organizations in more detail later in this book. The short version: treat them with caution. They exist, they serve a purpose in protecting draft specifications, but they are not a substitute for the discloser exercising judgment about what to share in a room full of competitors. Discloser beware.

One thing worth flagging here, because it surprises people who haven't worked at the institutional level: a meaningful part of the confidentiality you see around draft standards isn't really about protecting contributors. It's about protecting the standards organization itself. SDOs compete with each other — for relevance, for participants, for the right to host the canonical work in a given technology area — and they do not want a competing organization to see their draft work in progress and use it as a head start to scoop them. That competitive dynamic is one of the quieter reasons draft access tends to be restricted to members. It's worth keeping in mind, because it explains some confidentiality rules that would otherwise look excessive.

3.5 Confidentiality and the Open Source Tension

Many standards bodies restrict access to draft specifications to members only — though not all. Organizations like W3C and IETF make their drafts publicly available as a matter of policy, and the trend is toward more openness. But the members-only model remains common, particularly among older organizations and those with dues-dependent business models. There are two reasons for it, one financial and one strategic.

The financial reason is straightforward: membership dues are the primary revenue source for many organizations. If the drafts are available to everyone, the incentive to join diminishes.

The strategic reason is more interesting. Members who participate in the drafting process get an implementation head start. They're in the room. They can adapt the spec to their product needs. They can start building before the spec is published. By the time the final standard goes public, participating members may have a significant lead over non-members.

This creates a tension when standards bodies want to include open source implementations alongside their specs. Open source, by its nature, lives in public repositories. But if the spec is confidential, how do you have a public codebase implementing a private document?

Some organizations have tried to square this circle by keeping the open source code inside the organization, subject to the same confidentiality rules as the spec. The code uses open source licenses — all the rights are there — but you can't actually access it unless you're a member.

Others have simply accepted that once you commit to having an open source implementation, the spec is effectively public regardless of what the confidentiality rules say. A careful reader can reverse-engineer a lot from working code.

The broader pattern here is that as standards and open source converge, the traditional confidentiality model comes under pressure. Organizations that embrace open development tend to thrive. Those that try to keep one foot in each world tend to create confusion for their participants and friction for their lawyers.

Confidentiality in standards development should be treated with a light touch. It has a role — protecting draft specifications from premature publication, giving participants space to work without external pressure. But it should never be mistaken for a mechanism to protect proprietary information. That responsibility falls on the party doing the disclosing, not on the organization's confidentiality rules.

3.6 How the Pillars Interact

Copyright, patents, trademarks, and trade secrets don't operate in isolation in standards. They interact in ways that matter for how you advise.

The spec is protected by copyright — which maintains its integrity and prevents breaking forks. The technology described in the spec may be covered by patents. The brand and certification program are protected by trademarks. And trade secrets, while inappropriate for contributed technology, still require vigilance about what gets shared in rooms full of competitors. An implementer needs to navigate these layers to get a product to market.

In a well-designed standards ecosystem, these layers are complementary. The copyright ensures a canonical, stable spec that can't be fragmented. The patent policy ensures implementers can access the technology at reasonable cost (or for free, in RF regimes). The trademark program ensures interoperability and gives consumers confidence. And clear expectations about trade secrets and confidentiality keep participants from inadvertently compromising their own proprietary positions. Each layer solves a different problem.

Where things go wrong is when the layers conflict. When a certification program tests for patent-covered features that the spec describes as optional. When confidentiality rules prevent open source contributors from accessing the spec they're implementing. When a standards body's copyright enforcement posture discourages adoption of the standard it's trying to promote.

Most of the time, these conflicts are manageable. But recognizing them early is part of the job. When you pick up a new standards engagement, map all the IP dimensions before you dive into the patent policy. The patent policy is the hardest part, but it's not the only part.

We'll spend the next six chapters inside it.

Practice Tips

1. Read both inbound and outbound copyright terms — asymmetric terms can restrict implementer access even when participation is open.
2. Understand that a RAND commitment is a commitment to license, not an actual license — the terms are negotiated bilaterally outside the standards body.
3. Flag any policy that is silent on optional spec portions — this can leave implementers of optional features unprotected.
4. If a certification logo is required for market access, recognize that you've effectively admitted infringement of any necessary claim by putting the logo on the product.
5. Do not rely on confidentiality rules to protect trade secrets — the onus is on you to control what you share.
6. Check normative references in every spec — they pull in the referenced spec's IP terms.

Chapter 4 — Due Process, Best Practices, and the Antitrust Connection

Standards organizations are a creature of antitrust law. This isn't a metaphor — it's a statement about the legal foundation on which the entire system rests.

Think about what a standards body actually is. A group of competing companies gets together, agrees on a single way to do something, and then collectively adopts that approach — often to the exclusion of alternatives. In almost any other context, competitors agreeing to do the same thing in the same way would raise serious antitrust concerns. A group of competitors agreeing on price would be per se illegal. A group of competitors agreeing to exclude a competitor's technology from the market would be actionable.

Standards get a pass because they are generally considered pro-competitive. They enable interoperability, reduce costs, create markets, and benefit consumers. But that pass is conditional. It depends on the process being open, fair, and not dominated by any single participant or interest group. The moment the process becomes a vehicle for one company to impose its technology on the market, or for a group of companies to exclude a competitor, the antitrust protection erodes.

This chapter covers what makes a standard "open," the due process criteria that drive government recognition of private-sector standards, the antitrust protection that flows from those same criteria, the practical guardrails that standards bodies use to stay within bounds, and the confidentiality provisions that create tension with the openness principle.

4.1 What Makes a Standard "Open"

The term "open standard" is used liberally and means different things to different people. It's worth separating three distinct senses that often get collapsed together:

1. **Open process** — the development process is open to participation, balanced, consensus-based, and provides a right to appeal. This is the "open" that matters for government recognition and antitrust protection, and it's the subject of the rest of this chapter.
2. **Open output** — the final specification is publicly available, often without charge. Some organizations require this; others put specs behind a paywall or restrict them to members. A standard can have an open process and a closed output, or vice versa.
3. **Open implementation terms** — the resulting technology can be implemented royalty-free (or on RAND terms with no fee, or under some other open licensing regime). This is a patent-policy question, covered in Part II.

These are related but distinct attributes, and a standard can have some without having others. An ANSI-accredited standard behind a paywall, developed through a royalty-bearing RAND policy, is still "open" in the process sense. A specification published free on GitHub, developed by a closed group of self-selected vendors, is not — regardless of how accessible the final document is.

From a governance and antitrust perspective, which is where this chapter sits, "open" primarily refers to the **process** — not the output, and not the licensing terms. The rest of this chapter uses "open" in that sense.

The ANSI Essential Requirements

The **ANSI Essential Requirements** define what "due process" means in the standards context. While most technology standards work happens in consortia that aren't ANSI-accredited, the Essential Requirements represent a baseline that influences governance design across the industry — even in organizations that don't formally follow them.

The core principles:

Openness. Participation must be available to materially affected parties, without undue financial or organizational barriers. This doesn't mean every standards body has to be free to join; most consortia charge dues. It means the membership criteria can't be designed to exclude legitimate stakeholders.

Lack of dominance. No single interest category, individual, or organization may dominate the process. This doesn't mean everyone has equal influence — it means no one can exert undue authority that excludes fair consideration of others' viewpoints.

Balance. The process should seek a balance of interests. If the participants are heavily weighted toward one industry segment or one set of interests, the organization should make proactive efforts to recruit broader representation.

Consensus. ANSI defines consensus as "substantial agreement" — not unanimity. There must be a process for considering all views, resolving objections, and documenting how objections were addressed.

Right to appeal. Written procedures must include an accessible appeals mechanism. A participant who believes the process was unfair or that their objection was improperly dismissed must have a path to challenge the decision through an impartial review.

Notification and transparency. There must be timely notice of standards development activities, with adequate opportunity for participation, comment, and deliberation.

Consideration of views and objections. All comments and objections — whether from members or through public review — must be documented, considered, and resolved. You can't simply ignore an objection; you have to address it, even if the resolution is to proceed over the objection with documented rationale.

Even if your organization isn't ANSI-accredited, these principles provide a useful checklist for governance design. Can materially affected parties participate on reasonable terms? Is there balance among the participants? Is there a documented process for resolving objections? Is there a right to appeal? If the answers are yes, your governance is on solid ground.

These criteria do two substantive kinds of work that the rest of this chapter takes up in turn. First, they determine whether a standard qualifies for government recognition — the subject of the next section. Second, they form the primary defense against antitrust challenge — discussed in the section after that. The overlap is not coincidental: the same government concern about unfair competitive processes drives both bodies of rules. But the tests aren't identical, and a standard that satisfies one framework may still face scrutiny under the other.

4.2 Why the Criteria Matter: Government Recognition

The due process principles above aren't just good governance. They're the criteria that determine whether a standard gets picked up by government — for procurement, for incorporation into regulation, and for the presumption of legitimacy that comes with official recognition. Three frameworks are worth knowing by name.

United States — NTTAA and OMB Circular A-119. The National Technology Transfer and Advancement Act of 1995 directs federal agencies to use voluntary consensus standards developed by private-sector bodies in lieu of government-unique standards where feasible. **OMB Circular A-119** operationalizes that policy. A-119 defines a "voluntary consensus standards body" using criteria that map closely to the ANSI Essential Requirements — openness, balance, due process, consensus, appeals — because the two frameworks co-evolved. The practical consequences are meaningful: A-119-compliant standards are preferred in federal procurement, can be incorporated by reference into federal regulations (which is how a lot of U.S. technical regulation actually works), and carry presumptive legitimacy if challenged.

European Union — Regulation (EU) No 1025/2012. The European counterpart recognizes three European Standardisation Organisations (CEN, CENELEC, ETSI) as the bodies whose outputs can serve as harmonised European standards, and it includes openness and balanced-participation requirements similar to A-119's. The distinctive element is the **presumption of conformity** under the New Legislative Framework: when an EU directive sets essential requirements for a product and the Commission issues a standardisation request, a resulting harmonised standard cited in the Official Journal gives implementers a legal presumption of compliance with the underlying requirement. Compliance with the standard equals compliance with the law. This makes harmonised European standards considerably more legally significant than most U.S. voluntary consensus standards, and it has generated real litigation — including recent European Court of Justice decisions on public access to harmonised standards.

International — WTO TBT. Sitting above both frameworks is the WTO Agreement on Technical Barriers to Trade, which obligates members to use relevant international standards as the basis for technical regulations where feasible. The TBT Committee's 2000 Decision on Principles for the Development of International Standards identifies six principles — transparency, openness, impartiality and consensus, effectiveness and relevance, coherence, and attention to developing-country needs — that are the shared DNA of A-119, Regulation 1025/2012, and most serious consortium governance frameworks.

Two practical takeaways before we turn to antitrust:

The frameworks converge more than they diverge. If your governance satisfies one, it's likely to satisfy the others. Design for the highest common denominator.

The stakes are highest in the EU. Because harmonised European standards carry a presumption of conformity with EU legal requirements, participation in CEN, CENELEC, or ETSI work that supports a standardisation request involves more than technical influence — it involves shaping what compliance with EU law actually looks like. Counsel should recognize when a standards engagement has crossed from voluntary technical coordination into quasi-regulatory territory, because the analysis changes.

4.3 Why the Criteria Matter: Antitrust Protection

Government recognition is one reason the due process criteria matter. Antitrust protection is the other — and for most practitioners, it's the more immediate concern. The governance rules, voting mechanics, and decision-making procedures discussed in Chapters 11 and 12 aren't just organizational housekeeping. They help address potential antitrust risk by ensuring the process remains open, fair, and well-documented.

Why Following Process Matters Legally

Standards developed through a process that meets due process requirements carry greater legitimacy — in government procurement, in regulatory contexts, and in antitrust analysis. A standard developed through an open, balanced, consensus-based process with a right to appeal is far less likely to face an antitrust challenge than one developed behind closed doors by a self-selected group.

Conversely, a standard developed through a process that lacks openness, that is dominated by a single company, or that ignores dissenting views creates antitrust exposure for the participants. If a competitor argues that the standard was used as a tool to exclude their technology from the market, the quality of the process becomes the primary defense.

Every governance design choice — how votes are structured, who gets to participate, how objections are handled, what thresholds apply — should be evaluated not just for operational efficiency but for antitrust soundness. The two are related: a process that is fair, open, and well-documented is both operationally sound and legally defensible. A process that cuts corners on due process may be faster in the short term but creates risk that can undermine the entire standard.

The Voluntary Nature of Standards

One principle that underpins the antitrust analysis is that standards are voluntary. Nobody is required to implement a standard. Nobody is required to participate in its development. The market decides whether to adopt it.

This voluntariness is part of what makes standards pro-competitive. If a standard doesn't serve the market, the market ignores it. If a participant doesn't like the direction, they can leave and pursue alternatives. The threat of non-adoption disciplines the process — because a standard that reflects only one company's interests will be ignored by everyone else.

The voluntariness principle also limits the antitrust exposure of most standards activities. A group of competitors agreeing on a technical specification that anyone can implement is different from a group of competitors agreeing on pricing or market allocation. The former creates value for the market. The latter extracts it.

But voluntariness has limits as a defense. When a standard achieves monopoly-level adoption — when it becomes effectively mandatory to participate in a market — the voluntary label becomes less convincing. This is where patent policy, discussed in Chapters 7 through 10, becomes critical. If you must implement the standard to participate in the market, and the standard requires patents that a single company controls, the voluntary nature of the standard doesn't protect the market from monopoly behavior.

4.4 Antitrust Policy in Standards Bodies

The Guardrails: What You Can and Cannot Discuss

Most standards bodies have an antitrust policy that participants must acknowledge — often at the beginning of every meeting. The typical policy reminds participants that they are competitors and that certain topics are off limits.

You cannot discuss pricing, market allocation, customer allocation, production levels, bidding strategies, or any agreement to boycott a competitor or its products. You cannot use the standards process to collectively agree to exclude a competitor's technology for competitive reasons (as opposed to technical merit). You cannot use the standards body as a forum for coordinating commercial behavior.

You can discuss technical specifications, interoperability requirements, testing methodologies, conformance criteria, and the other substantive topics that standards bodies exist to address.

The line between permissible and impermissible discussion can be subtle. A conversation about which technical approach the standard should adopt is legitimate. A conversation about why a competitor's approach should be excluded because it would give them a market advantage is not — even if the technical discussion and the competitive discussion lead to the same outcome. The intent and framing matter.

Competitor Collaboration and Its Limits

Standards work is inherently collaborative among competitors. The antitrust framework allows this collaboration because it produces pro-competitive outcomes — interoperability, reduced costs, consumer choice. But the collaboration must stay within bounds.

The most common risk areas:

Patent disclosure and licensing discussions. Patent disclosures happen within the standards body — that's the point of the disclosure process. If you have a patent that may read on the standard, you disclose it to the working group. But licensing discussions — the actual negotiation of terms, rates, and conditions — must happen bilaterally, outside the standards body, between the patent holder and the implementer. The standards body does not participate in, broker, or administer these licensing discussions. A group of competitors collectively deciding to work around a particular company's patents — even if those patents are legitimately disclosed — can look like a group boycott. Individual companies can independently decide to design around a patent. They cannot coordinate that decision within the standards body.

Membership and participation restrictions. Excluding a competitor from a standards body, or from a particular membership tier, requires objective justification. If the exclusion looks like it's designed to disadvantage a competitor rather than to serve the standards process, it creates antitrust risk.

Certification and compliance programs. If a standards body controls a certification mark that is effectively required for market access (like the Wi-Fi logo), the certification criteria must be objective and non-discriminatory. Using the certification process to disadvantage a competitor's products — for example, by testing for features that only one company's products support — is problematic.

Standard-Setting as a Potential Antitrust Weapon

The standards process can be used offensively. A group of companies can form a standard around their preferred technology, excluding a competitor's approach, and then use market adoption of the standard to disadvantage the competitor. This is sometimes called "standards capture."

The defense against standards capture is the same as the defense against other antitrust risks: open process, broad participation, objective technical criteria for decisions, and documentation. If the decision to adopt one technical approach over another was made through an open process based on technical merit, the fact that the outcome disadvantages a competitor is not an antitrust violation. Competition has winners and losers. The process has to be fair; the outcome doesn't have to be equal.

Variations in Antitrust Policy Across Organizations

Antitrust policies vary significantly across standards bodies. Some have detailed policies with specific examples of permissible and impermissible conduct. Others have a one-paragraph reminder that gets read aloud at the beginning of each meeting.

International organizations may need to account for antitrust laws in multiple jurisdictions — US, EU, and other regimes that have different emphases and different enforcement patterns.

My view on antitrust policies in standards bodies: they're like "do not rob a bank" policies. Regardless of what the policy says — or doesn't say — you still can't violate antitrust law. A detailed antitrust policy doesn't create obligations that wouldn't exist without it, and a minimal policy doesn't create permissions that the law doesn't provide. The law applies whether the policy mentions it or not.

That said, a good antitrust policy serves a practical function. It reminds participants — many of whom are engineers, not lawyers — that they're sitting in a room with competitors and that certain topics are off limits. It creates a record that the organization takes compliance seriously. And it gives the chair a basis for redirecting conversations that stray into dangerous territory. The policy doesn't change the law, but it helps people follow it.

When engaging with a new organization, read the antitrust policy. Understand what it requires. And recognize that even in organizations with minimal written policies, the underlying antitrust principles still apply.

Antitrust policy isn't the only area where openness and pragmatism collide.

4.5 Confidentiality Provisions

Confidentiality in standards creates a tension with the openness principle, and it deserves careful handling.

When and Why Standards Work Is Confidential

Many standards bodies restrict access to draft specifications and working documents to members only. The reasons are both financial and strategic.

The financial reason: membership dues fund operations, and if drafts are freely available, the incentive to join diminishes.

The strategic reason: members who participate in the drafting process get an implementation head start. They can adapt the spec to their product needs and begin building before publication. By the time the final standard is public, participating members may have a significant lead over non-members. This early-mover advantage is a legitimate benefit of participation and an incentive for companies to engage.

Balancing Transparency with Competitive Concerns

Confidentiality rules exist to protect the development process — to give participants space to work, to propose ideas without public commitment, and to iterate without external pressure. These are legitimate purposes.

But confidentiality should not be used to protect proprietary information. As discussed in Chapter 3, standards are built with and among competitors, and they're intended for public availability. The onus is on the discloser to control what they share in a room full of competitors. Confidentiality rules protect draft specs, not business secrets.

The tension becomes acute when standards bodies want open source implementations alongside their specs. Open source lives in public repositories. Confidential specs live behind membership walls. Squaring this circle — as discussed in Chapter 3 — requires accepting that once you commit to an open source implementation, the spec is effectively public regardless of what the confidentiality rules say.

The broader trend is toward more openness. Organizations that develop in the open — with public repos, public mailing lists, and public meeting minutes — tend to attract broader participation and faster adoption. Those that maintain strict confidentiality tend to be older organizations or those in industries where the competitive advantage of early access is particularly valuable.

For practitioners, the key question is whether the confidentiality rules serve the process or whether they've become a barrier to the organization's goals. If the spec will ultimately be public, the value of interim confidentiality is limited. If the development process benefits from a period of protected iteration, confidentiality has a role — but it should be transparent about its purpose and limited in scope.

Practice note: The best advice on confidentiality in standards is to be as open as possible. Openness builds trust, attracts participation, and reduces the surface area for accusations of backroom dealing. Every layer of confidentiality you add is a layer that potential participants and implementers have to navigate — and some won't bother. Default to open unless there's a specific, articulable reason not to.

Practice Tips

1. Verify that the process meets the ANSI Essential Requirements or equivalent: openness, lack of dominance, balance, consensus, right to appeal, transparency.
2. Licensing discussions happen bilaterally outside the standards body — never within it.
3. Patent disclosures happen within the standards body — this is different from licensing.
4. Antitrust policies don't create new law — you can't violate antitrust whether the policy mentions it or not.
5. Default to open development unless there's a specific reason not to.

Chapter 5 — Open Source and Standards: The Convergence

Standards and open source started as separate worlds with separate cultures, separate legal frameworks, and separate communities. Over the past two decades, they've converged to the point where the line between them is often invisible. A specification developed through a standards process may be implemented as an open source project. An open source project may become the de facto standard. And increasingly, the same organization hosts both — with the same participants, the same governance challenges, and the same IP questions.

The convergence isn't an accident, and it's worth naming why. Standards and open source are both, fundamentally, structured collaboration among parties who don't fully trust each other and don't share interests in any deep way. Both rely on a layer of norms and practices — voting rules, contribution mechanics, IPR commitments, codes of conduct, review processes — that exist to make the collaboration possible at all. And in both, the essential ingredient is a creative tension between self-interest and collective benefit. Participants show up because contributing to the common artifact serves their own ends; the artifact is useful because no one participant can capture it. When that tension is in balance, the system produces things no participant could have built alone. When it isn't — when self-interest dominates and the work tilts toward one party, or when collective sentiment dominates and the work stops serving anyone in particular — the collaboration breaks down. Most of the legal machinery in both worlds is, at root, about keeping that tension productive.

This chapter explores that convergence: where the two cultures differ, where they overlap, and what practitioners need to understand to advise effectively in a world where the answer to "is this a standard or an open source project?" is increasingly "yes."

One caveat up front, because it's the single most common mistake counsel see in this area: **an open source license is not a substitute for a standards patent commitment.** An open source license (and any accompanying contributor patent grant) covers the specific code or text that contributors submit. It does not cover the specification as a whole, and it does not cover independent implementations by parties who read the spec and build their own version. A standards patent commitment covers necessary claims across the entire specification regardless of who contributed what. That is a fundamentally different scope of protection. Projects that release a "spec" under Apache 2.0, MIT, or Creative Commons and treat that as patent coverage are one bad disclosure away from a nasty surprise. Chapter 10 works through a live example. The rest of this chapter assumes that distinction and builds on it.

5.1 Two Cultures Colliding

The Standards Mindset vs. the Open Source Mindset

Standards development and open source development solve related problems but come from different traditions and operate with different assumptions.

Standards are rooted in antitrust law. A group of competitors comes together, agrees on a common approach, and publishes a specification that anyone can implement. The process emphasizes due process, balance, consensus, and the right to appeal. The output is a document — a specification — that describes what conforming implementations must do. The IP framework is built around patent commitments to that document.

Open source is rooted in copyright. An individual or group publishes code under a license that grants broad rights to use, modify, and redistribute. The process varies widely — from benevolent dictator models to formal governance structures — but it centers on contribution, review, and the accumulation of working code. The IP framework is built around copyright licenses and, in some cases, contributor license agreements.

These different foundations create different instincts. Standards professionals think in terms of specifications, conformance, and patent commitments. Open source professionals think in terms of code, licenses, and community governance. Standards professionals worry about scope creep and patent exposure. Open source professionals worry about license compatibility and maintainer burnout.

When the two communities work together — which they increasingly must — these different instincts can create friction. A standards professional may insist on formal patent commitments before work begins. An open source developer may view that as unnecessary overhead that slows development. An open source community may want to iterate quickly and publish frequently. A standards professional may insist on review periods, public comment, and formal approval processes.

Neither instinct is wrong. They reflect different risk models. The practitioner's job is to understand both and find structures that satisfy the requirements of each.

How the Cultures Differ in Practice

The cultural differences go beyond legal frameworks. They show up in how meetings are run, how decisions are made, and what counts as progress.

In standards, progress is measured by the maturity of the document. Working draft. Committee draft. Public review draft. Final specification. Each stage has formal criteria and usually a vote or consensus call. The document is the deliverable.

In open source, progress is measured by the quality of the code. Does it compile? Does it pass tests? Has it been reviewed? Has it been merged? The code is the deliverable. Documentation is important but secondary — and often written after the code, not before.

In standards, influence comes from organizational affiliation and formal roles (chair, editor, voting member). In open source, influence comes from contribution — the people who write the most code, review the most pull requests, and respond to the most issues earn the most influence, regardless of who employs them.

In standards, confidentiality during development is common. Draft specs are often restricted to members. In open source, development happens in public by default. The code is visible. The discussions are visible. The review process is visible.

These differences create practical challenges when the same project tries to operate in both modes. An engineer who's comfortable with open source workflows may find the standards review process painfully slow. A standards professional who's used to member-only access may be uncomfortable with everything happening in a public GitHub repository.

Code as Specification

The traditional model is clear: the standard is a document, and the code is an implementation of that document. They're separate artifacts governed by separate legal frameworks.

That model is breaking down. In an increasing number of cases, the code is the specification. The open source project is the canonical reference for how the technology works. There may be documentation — API references, protocol descriptions, developer guides — but the authoritative source is the codebase itself.

This creates IP complications. A standards patent commitment covers the specification — the document that describes what implementations must do. If there is no document — only code — what does the patent commitment cover? An open source license covers the specific code. It doesn't cover independent implementations by parties who read the code and build their own version. That difference is exactly what matters when someone independently implements a protocol described only by open source code.

This is the distinction discussed in Chapter 10's analysis of open source licenses applied to specifications. It matters here because as more "standards" are developed code-first, more implementations exist without the patent protection that a traditional standards process would provide.

When Open Source Replaces Standards

In some cases, open source doesn't just complement standards — it replaces them entirely. The technology becomes ubiquitous not because it was standardized through a formal process but because the open source implementation was so dominant that alternatives became impractical.

Chromium is a current example. The web standards that govern how browsers work are formally developed at W3C and WHATWG. But the dominant open source browser engine — Chromium, which powers Chrome, Edge, Opera, Brave, and others — is the practical arbiter of what the web actually is. If a website works in Chromium, it works for the vast majority of users. If it doesn't, the website changes, not the engine. The formal standard describes the ideal. The dominant implementation determines the reality.

This pattern creates a different kind of lock-in than proprietary formats. It's not vendor lock-in in the traditional sense — the code is open, anyone can fork it, and the licenses allow free use. But when a single implementation is so dominant that it defines the platform, the distinction between "open source" and "de facto standard" becomes academic. And the governance of that implementation — who accepts pull requests, who sets the roadmap, who decides what goes in and what stays out — becomes the governance of the standard, whether or not anyone calls it that.

For practitioners, this means that advising on "standards" sometimes means advising on dominant open source projects. The IP framework is different (code licenses rather than patent commitments). The governance is different (maintainer decisions rather than consensus votes). The risk profile is different (fork risk rather than patent assertion risk). But the strategic questions — who controls the technology, who benefits from adoption, who bears the cost of change — are the same.

Sometimes the line is difficult to draw. Linux is code — an open source project governed by maintainers under the GPL. But Linux has very stable APIs that effectively function as standards. Applications are written to those APIs. Entire ecosystems depend on their stability. Nobody went through a formal standards process to define them, but they carry the same practical weight as any published specification. Is the Linux kernel API an open source project or a de facto standard? The honest answer is both — and the distinction matters less than understanding what your client is building to and what governs it.

5.2 Open Source as Competitive Strategy in Standards

Reference Implementations in Open Source

As discussed in Chapter 13, reference implementations were traditionally created by standards bodies to verify that the specification could be implemented. They were not intended for production use.

The shift toward open source reference implementations changes the dynamic. When the reference implementation is an open source project intended for production deployment, it becomes the path of least resistance for adoption. Why implement the spec from scratch when you can use the existing, tested, production-ready code?

This creates a powerful competitive dynamic. The company or community that controls the reference implementation has outsized influence on adoption — because most implementers will use their code rather than building their own. If the open source implementation becomes dominant, it becomes the de facto standard regardless of what the specification document says.

The WHATWG/W3C HTML5 Fork

The HTML5 story, discussed in Chapter 1, is the canonical example of open source and standards convergence. When W3C's direction for the web diverged from what browser vendors wanted, Apple, Mozilla, and Opera formed WHATWG and forked the HTML spec. Google and Microsoft joined later. The actual HTML development happens at WHATWG; W3C publishes it.

The lesson isn't just about organizational politics. It's about the power of implementation. WHATWG succeeded because its members controlled the browsers. They could implement their version of the spec directly. W3C's version — without browser implementation — was academic. The implementers determined the standard, not the standards body.

This pattern repeats across the industry. When the people who implement the technology are the same people who develop the specification — or when a dominant open source implementation exists — the formal standards process becomes a ratification of what the code already does, not a design exercise.

5.3 Hybrid Models

Community Specification License

The Community Specification License, discussed in Chapter 9, represents one approach to bridging the gap. It applies standards-style IP terms — patent commitments, exclusion mechanisms, contribution-triggered licensing — to work done in Git repositories using open source-style workflows.

The CSL acknowledges that modern specification development often looks more like open source development than traditional standards work. Small teams iterate on a spec in a repo, accept contributions via pull requests, and publish when ready. The CSL provides the patent clarity that open source licenses don't, without requiring the organizational overhead of a formal standards body.

Open Source Foundations as Standards Homes

The Linux Foundation, Eclipse Foundation, Apache Foundation, and others increasingly host both open source projects and standards initiatives. The organizational infrastructure — governance frameworks, IP policies, operational support — serves both purposes.

This co-location is natural. The communities overlap. The governance challenges are similar. And having both the spec and the code under the same organizational umbrella simplifies the IP relationship between them.

But co-location creates its own challenges. Open source governance (meritocratic, contribution-based, often informal) and standards governance (due process-based, consensus-driven, often formal) don't always coexist easily. A project that tries to operate under both models simultaneously can end up satisfying neither — too slow for the open source community, too informal for the standards community.

The organizations that handle this best are the ones that clearly separate the two functions — spec development under standards governance with appropriate patent policies, code development under open source governance with appropriate licenses — while ensuring the two tracks are coordinated. JDF's model, where the spec is governed by standards IP terms and the code is governed by an open source license, is one approach that works.

The key design question is institutional. Does the foundation treat the spec as an appendage of the code project? Or does it treat the spec as a first-class artifact with its own governance track? Organizations that default to open source governance for everything tend to produce specs with inadequate patent coverage. Organizations that default to standards governance for everything tend to produce code with processes too heavy for developer communities. The goal is to match the governance model to the artifact.

The Convergence in Practice: What It Looks Like Today

In practice, the convergence takes several common forms:

Spec-first with open source reference implementation. The traditional model, updated. A working group develops a specification through a standards process with patent commitments. An open source implementation is developed alongside — sometimes by the same people, sometimes by a separate community. The spec is authoritative; the code validates it. This model is used by many JDF projects and OASIS technical committees.

Code-first with spec documentation. An open source project develops a technology, gains adoption, and then documents it as a specification — either for interoperability (so others can build independent implementations) or for formalization (to submit through a PAS process for international recognition). The code is authoritative; the spec describes it. This is increasingly common for API definitions and protocol work.

Dual-track parallel development. The spec and the code are developed simultaneously by overlapping teams. Neither is strictly authoritative — they inform each other iteratively. This is the hardest model to manage but can produce the best results when the coordination is deliberate.

Open source as de facto standard with no spec. The code is the standard. There is no separate specification document. Implementers who want to interoperate read the code, reverse-engineer the behavior, or use the code directly. This is functionally what happens with Chromium for web rendering, and it's what the Crystal Ball chapter (Chapter 17) suggests may become more common with AI-driven interoperability.

AI and the New Frontier

AI is accelerating the convergence. The Agentic AI Foundation at the Linux Foundation, Google's Agent-to-Agent Protocol, and similar initiatives are developing interoperability specifications for AI systems. Some use traditional standards approaches. Others use open source approaches. Most use some hybrid.

As discussed in Chapter 17, the question of which approach is appropriate depends on the artifact. Protocol specifications benefit from standards-style patent commitments. Code benefits from open source licenses. Model artifacts, training data documentation, and evaluation frameworks don't fit neatly into either category.

This is the frontier for standards and open source convergence. The frameworks discussed in this book — patent policies, governance structures, non-assert agreements, the Community Specification License — provide the tools. Adapting them to new types of artifacts is the work ahead.

5.4 Practical Considerations for Attorneys

License Selection for Spec-Adjacent Code

When choosing a license for code that implements or accompanies a specification, the interaction between the open source license and the standards patent policy matters.

Permissive licenses (Apache 2.0, MIT) are most commonly used alongside standards work. Apache 2.0 includes an explicit patent grant tied to contributions, which provides some patent coverage — though narrower than a standards patent commitment. MIT includes no explicit patent grant at all.

Copyleft licenses (GPL, LGPL) are rarely used in standards-adjacent contexts. The copyleft obligation to share modifications can conflict with the proprietary implementation model that many standards participants rely on. In practice, copyleft licenses are essentially never used by standards bodies for their own reference implementations.

The key guidance: don't assume the open source license provides the same patent coverage as the standards patent policy. They cover different things. If both the spec and the code exist, make sure your client understands which rights come from which framework.

Governance Alignment

When a project has both a spec track and a code track, the governance structures need to be coordinated. Who has authority over the spec? Who has authority over the code? What happens when the code diverges from the spec? What happens when the spec evolves but the code doesn't follow?

A threshold question that must be answered at the outset: **which leads?** Is the spec being drafted first, with the code implemented against it as a reference? Or is the code leading, with the spec documented against a code snapshot at a point in time? These are fundamentally different development models with different IP implications.

In a **spec-leads model**, the specification is the authoritative source. The code implements what the spec describes. Patent commitments attach to the spec, and the code is a derivative artifact. Changes flow from the spec to the code, and conformance is measured against the document.

In a **code-leads model**, the codebase is the authoritative source. The spec documents what the code does. The spec may lag the code, and the code may evolve between spec snapshots. Patent commitments, if they attach to the spec, may not cover the latest code — and if they attach to the code, the framework described in Chapter 10 (open source licenses providing narrower patent coverage than standards commitments) applies.

Many projects fall somewhere in between, with the spec and code evolving in parallel. This can work if the coordination is deliberate — but if no one has decided which leads, the two tracks will diverge, and the divergence will surface at the moment someone asks "which version is authoritative?"

In traditional standards organizations, code was used primarily to verify the specification — reference implementations that proved the spec could be implemented and that identified ambiguities or errors. These reference implementations were not intended for production use, were often unoptimized, and frequently carried restrictive licensing terms specifically to prevent anyone from deploying them as products. The intent was to keep the spec canonical and the code subordinate. The modern shift toward production-quality open source implementations alongside specs inverts that relationship in ways that the traditional governance model wasn't designed to handle.

These questions are easier to answer when they're addressed at the design stage. Define which governance applies to which artifact. Define how conflicts between the spec and the code are resolved. Define who has the authority to declare that the code is a conformant implementation.

If you don't define these boundaries, the open source community and the standards community will each assume they have authority — and the resulting governance confusion will consume more time and energy than the underlying technical work.

Practice Tips

1. Don't assume the open source license provides the same patent coverage as the standards patent policy.
2. When a project has both a spec and code, decide at the outset which leads.
3. Separate governance tracks for spec and code — match the governance model to the artifact.
4. If no specification document exists and the code is the standard, understand the patent coverage gap for independent implementers.

Chapter 6 — Necessary Claims: The Heart of Patent Policy

Every patent policy in every standards body revolves around one concept: **necessary claims**. Sometimes called essential claims, these are the patent claims that an implementer cannot avoid when building to a standard. If there's an alternative way to achieve the same result without practicing the claim, it's not a necessary claim — because competition among alternatives keeps pricing and licensing terms in check. But if there is no alternative, the patent holder has a monopoly created by the standard itself, and without the constraints of a patent policy, could abuse that position through unreasonable pricing, onerous terms, or injunctions.

This is the problem that RAND and royalty-free commitments solve. They restore the results of competition — reasonable pricing, accessible terms — in situations where the standard has effectively removed competition. The necessary claims definition is the mechanism that determines when those commitments apply.

6.1 The Logic: Blocking Patents and Market Discipline

In a normal market, competition among alternative technologies imposes discipline on patent holders. If someone holds a patent on one approach, an implementer can design around it, license a competitor's alternative, or build something different. The availability of alternatives keeps pricing and licensing terms reasonable — not because the patent holder is generous, but because the implementer has options.

Standards remove those options. When a technology is embedded in a standard, every implementer must use it. There is no designing around it — the standard says do it this way. A patent holder whose claims read on "this way" could, without a patent policy commitment, abuse that monopoly position: charge whatever they want, impose onerous terms, seek injunctions, or simply refuse to license. The implementer has no alternative and no leverage.

RAND and royalty-free commitments act to restore the results of competition. Even though the standard has eliminated competing approaches, the patent policy requires the patent holder to behave as if competition still existed — offering licenses on reasonable terms (in RAND) or at no charge (in royalty-free). The necessary claims definition determines when this obligation kicks in: only for claims that are truly unavoidable. Where alternatives exist, market forces do the work on their own.

This is why standards focus on the **what**, not the **how**. A specification describes the interface — the format of the data, the protocol for communication, the expected behavior. It generally does not prescribe how the implementation achieves those results internally. A number-sorting standard might say: accept a series of numbers, return them sorted lowest to highest. It doesn't say use bubble sort or quicksort. A patent on bubble sort wouldn't be a necessary claim because the implementer can choose a different algorithm. A patent on the act of sorting numbers in the format specified by the standard — if such a patent existed and couldn't be avoided — would be.

6.2 Parsing the Definition

The typical necessary claims definition reads something like: "claims of patents or patent applications, owned or controlled by a member or its affiliates, now or at any time in the future, that would necessarily be infringed by an implementation of the specification."

"Claims of patents or patent applications"

The commitment covers both issued patents and pending applications. This is intentional — if it only covered issued patents, a participant could keep their application pending throughout the standards process and only assert once the patent issues, after the standard is locked in and widely deployed. Including applications closes that gap.

"Owned or controlled"

The commitment extends beyond patents the participant directly owns to patents they control — through exclusive licenses, governance rights over an affiliate, or other legal arrangements. The intent is to prevent participants from structuring their patent holdings to avoid the commitment while retaining the economic benefit.

This generates real disputes. A venture-backed startup may be controlled by a VC that also controls hundreds of unrelated companies. A conglomerate's participation through one subsidiary doesn't necessarily mean it can bind patents held by a completely separate division. These are legitimate structural issues with no universal solution. Sometimes the answer is a side letter clarifying scope. Sometimes the answer is that the participant needs to sort out their internal arrangements before joining.

"Now or at any time in the future"

The commitment is forward-looking. If you acquire a patent after making the commitment — through R&D, purchase, or corporate acquisition — and that patent reads on the standard, it's pulled into the commitment. You can't make a commitment on Monday and buy a blocking patent on Tuesday.

"Necessarily be infringed by an implementation"

This is the core of the definition. The claim must be unavoidable. If there's any way to implement the specification without practicing the claim, it's not a necessary claim. The commitment doesn't cover the entire patent portfolio — just the claims that the standard forces implementers to practice.

6.3 Normative vs. Optional: The Empty Commitment Problem

Standards specifications typically contain both **normative** elements (things an implementer must do) and **optional** elements (things an implementer may do). The question of whether the necessary claims definition covers optional portions is more significant than it might appear.

If the definition only covers normative, required elements, then anything marked optional in the spec is outside the commitment. An implementer who chooses to implement an optional feature gets no patent protection for that feature. The patent holder can charge whatever they want — or seek an injunction — for optional elements.

This matters because many specifications are predominantly optional. The office file formats, for instance, are roughly 95% optional features. If you want to build a word processor that doesn't support bold text, the spec allows that. But nobody builds a word processor without bold text. If the patent commitment only covers the handful of truly mandatory elements, it may cover very little of what implementers actually need.

This criticism was leveled against some early patent grants in the industry — the argument being that because nearly everything was optional, the licensing commitment was effectively empty. The intent wasn't to create an empty commitment, but the drafting didn't anticipate the issue.

More modern policies address this explicitly. They typically cover the **normative elements of optional portions** — meaning that if you choose to implement an optional feature, the patent commitment extends to the required elements within that feature. You don't have to implement bold text, but if you do, the patent commitment covers it.

If you're reviewing a policy that's silent on optional portions, flag it. The silence may be deliberate or may be an oversight, but either way it creates ambiguity that can undermine the value of the commitment.

"Shall" vs. "May" — The Language That Drives It

In standards drafting, the words "shall" and "may" have specific, consequential meanings. "Shall" denotes a normative requirement — an implementer must do this. "May" denotes an optional feature — an implementer can do this but doesn't have to.

This distinction directly maps to the necessary claims analysis. A patent claim that reads on a "shall" element is more likely to be a necessary claim because every compliant implementation must practice it. A patent claim that reads on a "may" element is only necessary for implementers who choose to implement that option — and only if the policy extends the commitment to normative elements of optional portions.

When reviewing a specification alongside its patent policy, pay attention to how liberally the drafters used "shall" versus "may." A specification that marks most features as "may" reduces the mandatory surface area and, depending on the policy, may narrow the patent commitment significantly. This can be deliberate — a way for patent holders to limit their exposure — or it can be a natural consequence of a specification that genuinely offers implementers flexibility. Either way, the language choices in the spec and the necessary claims definition in the policy interact to determine what's actually covered.

6.4 When Commitments Lock In

In open source, patent commitments typically attach at the moment of contribution. You submit code, the license terms apply, and the commitment is made.

Standards work differently. The patent commitment generally locks in when the specification is **finalized and approved** by the working group — not when individual contributions are made during the drafting process. This means that during development, contributions are provisional. The commitment crystallizes at the end, and it covers the final deliverable as a whole, regardless of who contributed what.

This has several implications.

First, it means you're making a commitment to the entire final specification, not just to the parts you contributed. If another participant introduces technology that reads on your patents, your commitment covers it — assuming it ends up in the final spec and your claims are necessary.

Second, it creates a window during development where participants can evaluate the evolving spec and decide whether to stay or withdraw. If the spec moves in a direction you didn't anticipate — you thought you were building a microwave standard, and the group decided to make a toaster — you have the opportunity to withdraw before the commitment locks in. This microwave-becomes-a-toaster problem is one of the most common sources of friction in standards work, and it's why scope — the definition of what a working group can and cannot work on — matters so much. We'll discuss scope in depth in Chapter 8.

Third, it means that interim drafts, working documents, and rejected proposals generally don't carry patent commitments. Only the approved final deliverable does.

6.5 The Contribution Trigger vs. the Participation Trigger

Patent policies use one of two triggers for the commitment: participation or contribution.

A **participation trigger** means that by joining the working group, you've made the commitment. Your necessary claims are subject to the licensing terms for whatever the group produces, regardless of whether you contributed any technology. You're in the room, you're bound.

A **contribution trigger** means the commitment only attaches to technology you actually contributed. If you participate as an observer and contribute nothing, you have no patent commitment. But the commitment on what you do contribute may be broader — sometimes covering not just the specific contribution but anything in the final spec that relates to it.

The participation trigger is more protective for implementers. It ensures that everyone in the room has skin in the game, and it prevents a participant from sitting quietly while their patents get embedded in the spec and then asserting without constraint.

The contribution trigger is more protective for patent holders. It limits exposure to what you actually brought to the table. The tradeoff is that it creates a category of participants — observers with relevant patents — who benefit from the standard without contributing to the patent commitment pool.

Most modern royalty-free policies use a participation trigger. RAND policies are more varied. Either way, understanding which trigger your policy uses is essential for advising clients on what joining a working group actually means.

6.6 Successors and Assigns

What happens to the patent commitment when patents change hands? If a participant sells a patent that's subject to a standards commitment, does the commitment travel with the patent?

Most well-drafted policies address this by requiring that any transfer of committed patents be subject to the existing commitment. The buyer takes the patent subject to the licensing obligation. This prevents a participant from making a commitment and then selling the patent to a third party that asserts freely.

In practice, enforcement of successor provisions is uncertain. A patent troll that acquires a patent may argue it had no knowledge of the standards commitment, or that the commitment doesn't run with the patent as a matter of law. Courts haven't fully resolved these questions, and the answers may vary by jurisdiction.

What this means practically is that successor and assign provisions are necessary but not sufficient. They provide a contractual basis for arguing the commitment survives a transfer, but they may not prevent assertion by a determined acquirer. This is one of the areas where the theoretical framework of patent policies meets the messy reality of patent transactions.

6.7 The Versioning Question

One of the most debated areas in necessary claims practice is how commitments carry across specification versions.

You participate in version 1.0 of a specification. Your necessary claims are committed. Then the group begins work on version 2.0. Several scenarios arise.

You stay and participate in 2.0. Your commitment extends to 2.0 through the normal operation of the policy — you're participating, the spec is finalized, the commitment locks in.

You leave after 1.0, and 2.0 adds new sections without changing 1.0. Your commitment clearly covers 1.0. It likely doesn't cover the new sections in 2.0, since you weren't involved. But the 1.0 material is still in the field, still being implemented, and your commitment to it survives your departure.

You leave after 1.0, and 2.0 modifies the 1.0 material. This is where it gets contested. Your commitment was to the 1.0 text. The 2.0 text is different, even if your patents still read on it. Does your commitment carry forward?

Some policies address this with a "substantially similar" formulation — if the technology is used in a substantially similar way for a substantially similar purpose, the commitment persists. Others treat each version as an entirely new commitment. Many are silent, which creates ambiguity that tends to surface at the least convenient time.

A related issue arises with **normative references**. If version 2.0 of a specification includes a normative reference to version 1.0 — essentially saying "go implement v1.0, then come back here for the new material" — the v1.0 commitment is preserved intact because the text hasn't changed. But if the v2.0 team copies the v1.0 text into the new document and edits it, the commitment to the original text may not clearly extend to the modified version.

These aren't theoretical problems. They arise in practice when participants withdraw from working groups and later seek to assert patents against the continued work. The best protection is clear policy language. The second-best protection is a normative reference structure that avoids modifying committed text.

6.8 Software and Implementation Patents

One area that trips up newcomers is the distinction between patents on the standard and patents on the implementation.

The necessary claims definition covers patents that read on the **specification** — the normative description of what an implementation must do. It generally does not cover patents on **how** the implementation achieves that result internally.

Going back to the sorting example: a patent on accepting numbers in the specified format and returning them sorted as described would be a necessary claim. A patent on a particular sorting algorithm used internally would not — because the standard doesn't prescribe the algorithm.

This distinction is what allows competition among implementers. Everyone implements the same interface, but they compete on how efficiently, how quickly, or how cleverly they do the internal work. Implementation patents remain outside the standards commitment and subject to normal market dynamics.

The boundary isn't always clean. Encryption standards, for instance, sometimes require specific algorithms — the standard can't be implemented without using the exact algorithm described. In those cases, patents on the algorithm itself may be necessary claims. But those are the exception. For most standards, the interface is what's committed; the implementation is what's competed.

6.9 The Game Theory of Necessary Claims

As discussed briefly in Chapter 3, there's a counterintuitive dynamic in how necessary claims are argued in disputes.

When a patent holder asserts a claim against an implementer, both sides generally have an incentive to agree the claim is a necessary claim — even though they'll fight over everything else.

The implementer wants the claim to be necessary because that brings it within the RAND or royalty-free commitment. The patent holder gets a reasonable royalty (in RAND) or no royalty (in RF), but the implementer avoids unconstrained damages and injunctive relief.

The patent holder also often prefers the claim to be necessary — paradoxically — because proving infringement is dramatically simpler. If the implementer has a certification logo on their product (Wi-Fi, Bluetooth, etc.), they've already told the world they implement the standard. The patent holder doesn't need to prove infringement claim by claim. The only question is the rate.

If the implementer successfully argues the claim is *not* necessary — just an implementation detail — they've escaped the standards commitment but walked into a regular patent infringement case.

No RAND protection. No rate cap. Full damages and injunctive relief are on the table.

This dynamic is why disputes over necessary claims are surprisingly rare compared to disputes over rates. Both sides prefer to be inside the system, arguing about the price, rather than outside it, arguing about liability.

One related and highly contested question — whether a RAND-committed patent holder can seek injunctions against implementers — is addressed in Chapter 7.

Practice Tips

1. Distinguish between patents on the specification (necessary claims) and patents on implementation details (not covered).
2. Identify whether the policy uses a participation trigger or contribution trigger.
3. Verify that the necessary claims definition covers normative elements of optional portions.
4. Pay attention to "shall" vs. "may" language in the spec — it directly determines the surface area of necessary claims.
5. Check whether the commitment covers patents acquired after the commitment is made.
6. Confirm how successor and assign provisions work — if patents change hands, the commitment should follow.

Chapter 7 — RAND and Royalty-Bearing Standards

RAND — Reasonable and Non-Discriminatory — is the patent policy model that allows patent holders to charge royalties for their necessary claims while constraining how much they can charge and on what terms. It is the dominant model in industries where companies have invested heavily in patented technology and expect to recover that investment through licensing: wireless communications, video codecs, cellular networks.

The goal, as discussed in Chapter 6, is to restore the results of competition. Even though a standard eliminates competing approaches, the RAND commitment requires the patent holder to behave as if alternatives still existed — licensing at market rates rather than monopoly rates. The patent holder can monetize their investment. The implementer gets access on reasonable terms. That's the theory.

Proponents of RAND argue that it's necessary to incentivize innovation — that companies won't invest in developing new technology for standards unless they can monetize the resulting patents. There is no empirical proof of this. The web was built almost entirely on royalty-free standards, and it's difficult to argue that the web lacks innovation because patent holders weren't collecting royalties on HTTP, HTML, or CSS. The reality is that RAND and royalty-free reflect different business models, not different levels of innovation. Some industries have established patent licensing as a core revenue stream. Others haven't. The choice of patent policy follows from that commercial reality.

This chapter covers how RAND works in practice, how disclosure obligations operate, how disputes over rates play out, and the contentious question of whether RAND-committed patent holders can seek injunctions.

7.1 When Organizations Choose RAND

The choice between royalty-free and RAND is one of the first and most consequential decisions in setting up a standards body.

Royalty-free works when the participants are primarily interested in broad adoption and commodity interoperability. The web standards world is a natural fit — everyone benefits from ubiquitous implementation, and the value is in the products and services built on top of the standard, not in the standard itself.

RAND works when some participants hold valuable patent portfolios developed through significant R&D investment, and they expect to license those patents as part of their business model. The telecom industry is the canonical example. Companies spend years developing wireless technology, file patents on it, and then contribute it to standards bodies with the expectation that they'll license those patents to implementers.

It's worth noting that "RAND" doesn't mean everyone charges royalties. A RAND commitment means you *can* charge, not that you *must*. The vast majority of RAND-committed patent holders never set up active licensing programs. They make the commitment, participate in the work, and never seek royalties. Active licensing is concentrated in a few specific industries.

7.2 FRAND: Fair, Reasonable, and Non-Discriminatory

In Europe, the commitment is typically called FRAND — with the F standing for Fair. RAND and FRAND mean the same thing in practice, and no, the F does not mean Free. This is a common misunderstanding that can be consequential.

A RAND or FRAND commitment is a promise to offer a license to necessary claims on reasonable and non-discriminatory terms, negotiated bilaterally between patent holder and implementer outside the standards body. As discussed in Chapter 3, this is a commitment to license — not a license itself. The standards organization facilitates the commitment but doesn't set rates, administer licenses, or broker negotiations.

What Does "Reasonable" Mean?

This is where all the litigation happens. The patent policy terms themselves are rarely contested in court. It's the rates that get fought over, and those fights can involve hundreds of millions of dollars.

"Reasonable" is generally understood to mean the rate the patent holder could command in a hypothetical negotiation where competitive alternatives exist — even though, in reality, the standard has eliminated alternatives. The idea is to strip out the monopoly premium created by standardization and arrive at the rate the patent would command on its own merits.

Courts have developed various methodologies for calculating RAND rates — comparable license analysis, top-down approaches based on the aggregate royalty burden for the standard, bottom-up approaches based on the patent's incremental value. The details are the domain of licensing litigators and economists and are beyond the scope of this book. What matters for practitioners working on the policy side is understanding that "reasonable" is a contested concept, that rates are negotiated or litigated outside the standards body, and that the policy creates the framework but not the answer.

What Does "Non-Discriminatory" Mean?

The non-discriminatory prong requires that the patent holder offer the same (or comparable) terms to similarly situated licensees. You can't charge one implementer a penny and another a dollar for the same technology without justification.

In practice, "non-discriminatory" doesn't mean identical terms for everyone. Volume discounts, cross-licensing arrangements, and differences in the licensee's use case can all justify different terms. What it prohibits is using the licensing process to pick winners and losers among implementers — charging a competitor more simply because they compete with you.

7.3 The Disclosure Obligation

RAND policies typically include a disclosure obligation: if you're aware of patents that may read on the standard, you must tell the group. This is distinct from the royalty-free model, where the focus is on exclusions. In RAND regimes, the focus is on transparency.

How Disclosure Works

Disclosure obligations come in the same two forms as exclusion mechanisms: a formal call for patents tied to a specific milestone in the specification process, or a rolling obligation to disclose whenever you become aware of a relevant patent.

The arguments for and against each approach parallel the exclusion discussion in Chapter 8. A formal call gives participants a concrete deadline and a trigger for internal review. A rolling obligation is self-executing and doesn't depend on the organization following its own procedures. For large organizations that may not have day-to-day visibility into what their standards participants are doing, the formal call is generally more effective — it produces a piece of paper that gets routed to the patent team.

The Standard of Knowledge

Most disclosure policies are based on the **actual personal knowledge** of the working group representative — not the collective knowledge of the entire company. You don't have to conduct a patent search. You don't have to consult every patent attorney in the building. If the individual who participates in the working group is personally aware of a potentially relevant patent, they must disclose it.

This standard is deliberately narrow. The rationale is that imposing a broader due diligence obligation would make participation too expensive and would deter companies from engaging. The tradeoff is that relevant patents can go undisclosed when the representative genuinely doesn't know about them — even if someone else in the organization does.

The narrow standard creates an obvious incentive problem: send someone who doesn't know your patent portfolio, and you can avoid disclosure. This strategy tends to cancel itself out in practice, because a company that fails to disclose and later tries to assert faces arguments about patent misuse, unclean hands, and — as one court found — antitrust liability.

The Qualcomm v. Broadcom Lesson

The most cited case on disclosure obligations is *Broadcom Corp. v. Qualcomm Inc.* Qualcomm participated in a standards body, failed to disclose patents it held on the standard, and later asserted those patents against implementers. The litigation was compounded by discovery misconduct — Qualcomm produced critical documents only on the eve of trial.

The court found that Qualcomm's failure to disclose constituted anticompetitive behavior. But the reasoning went further than the patent policy's own terms. Rather than applying the policy's disclosure standard — actual personal knowledge of the representative — the court articulated a broader standard based on the "reasonable expectations" of the other participants. This effectively created a disclosure obligation beyond what the policy's text required.

The case is important but problematic. It was decided against a backdrop of egregious misconduct — the court was looking for a way to sanction Qualcomm's behavior, and the disclosure failure provided a vehicle. Whether the "reasonable expectations" standard would be applied to a good-faith participant who simply didn't know about a relevant patent is unclear. The case hasn't been widely followed, but it remains a cautionary example of what happens when disclosure goes wrong.

No Duty to Search — But Disclosure Is Not Optional

Almost universally, RAND policies do not impose an obligation to conduct a patent search. The actual personal knowledge standard is the floor and the ceiling. This is worth emphasizing because participants — especially engineers — sometimes assume they need to do more. They don't. And the organizations themselves shouldn't be conducting searches on behalf of participants.

What does need to be emphasized — because this is a point many participants get wrong — is that **disclosure itself is mandatory, not optional**. If you have a call for patents and you have a patent that you're aware of, you must disclose it. This is not a discretionary decision. It is a requirement of the policy.

Engineers and business teams sometimes resist disclosure because they worry it will look bad, send a negative signal to the community, or derail the work. Those concerns are understandable but irrelevant. If the policy requires disclosure based on actual personal knowledge, and the representative knows about a relevant patent, there is no choice to make. The obligation is clear.

The failure to disclose a known patent doesn't eliminate the licensing obligation — the RAND commitment still applies whether or not you disclosed. But non-disclosure creates significant enforcement problems if you later try to assert. It can be characterized as patent ambush, it invites antitrust scrutiny, and it undermines your credibility in the community. Disclosure may feel uncomfortable, but the alternative is worse.

7.4 The Injunction Question

One of the most contested issues in standards patent law is whether a holder of a standard-essential patent who has made a RAND commitment can seek an injunction against an implementer. The arguments on each side are substantial.

The Case Against Injunctions

If a patent holder has committed to license on reasonable and non-discriminatory terms, seeking an injunction — an order that would prevent the implementer from shipping their product — is arguably inconsistent with that commitment. The patent holder has already agreed to license. The only open question is the rate. An injunction in this context becomes a tool for extracting above-market royalties under the threat of shutting down the implementer's business. This is the **patent holdup** problem: the patent holder uses the threat of an injunction to force a settlement far above what a reasonable royalty would be.

The Case For Injunctions

Patent holders argue that without the possibility of an injunction, implementers have no incentive to negotiate in good faith. If the worst outcome for an implementer is paying a reasonable royalty — the same amount they'd pay if they had negotiated up front — there's no cost to refusing to negotiate and simply implementing without a license. The implementer captures the full value of the standard while the patent holder bears the cost of enforcement. This is the **implementer holdout** problem: the implementer delays, ignores licensing overtures, and forces the patent holder to litigate just to get paid what was already owed.

How Courts Have Addressed It

Courts in different jurisdictions have taken different approaches, and the law is not settled.

In the United States, the Supreme Court's 2006 decision in *eBay Inc. v. MercExchange* established that injunctions in patent cases are not automatic. Courts must apply a four-factor equitable test: irreparable injury, inadequacy of monetary remedies, balance of hardships, and the public interest. In the SEP context, U.S. courts have generally been reluctant to grant injunctions where the patent holder has made a FRAND commitment, reasoning that monetary damages — a reasonable royalty — are an adequate remedy. The Ninth Circuit reinforced this in *Microsoft Corp. v. Motorola Inc.*, finding that a demand for an injunction on FRAND-committed patents was inconsistent with the licensing obligation.

In Europe, the Court of Justice of the European Union took a different approach in *Huawei Technologies v. ZTE Corp.* (2015). The CJEU established a structured negotiation framework: before seeking an injunction, the SEP holder must notify the implementer and make a specific FRAND licensing offer. The implementer must respond diligently and in good faith. If the implementer refuses to engage — if they are an "unwilling licensee" — the patent holder may seek an injunction without running afoul of EU competition law. This framework gives both sides procedural obligations and reserves injunctions for cases where the implementer is genuinely refusing to negotiate.

The practical result is that injunctions on FRAND-committed patents are rare in the United States and available but procedurally constrained in Europe. Other jurisdictions — China, the UK, India — are developing their own approaches, and the interaction of multiple jurisdictions in global licensing disputes adds further complexity.

Why This Matters for Policy Design

For practitioners drafting or evaluating patent policies, the injunction question is relevant even though it plays out in courts rather than in the standards body itself. A policy that is explicitly silent on injunctions leaves the question to the courts. Some policies include language that the commitment is to license on RAND terms "without seeking injunctive relief," though the enforceability of such provisions varies. Others deliberately leave it open, taking the position that injunctive relief is a judicial question, not a contractual one.

The injunction debate is ultimately about the balance of power between patent holders and implementers — and that balance differs depending on the industry, the jurisdiction, and the specific facts. Detailed treatment of post-implementation licensing disputes is beyond the scope of this book, but understanding the injunction dynamic is essential context for the patent policy choices covered here and in Chapter 8.

7.5 Patent Pools and Collective Licensing

In industries where a standard is covered by patents from many holders — video codecs and wireless are the primary examples — **patent pools** often form to simplify licensing.

A patent pool aggregates patents from multiple holders into a single licensing program. Instead of negotiating separate licenses with dozens of patent holders, an implementer takes one license from the pool and gets access to all the pooled patents. The pool administrator collects royalties and distributes them to the patent holders according to an agreed formula.

Patent pools are generally formed outside the standards body, often by a licensing administrator like MPEG LA or Via

Licensing. The pools typically employ an independent expert who evaluates whether each submitted patent is truly a necessary claim before it's included in the pool. This gatekeeper function is important — without it, patent holders could submit marginal patents and dilute the pool's value.

One thing to watch: the independent evaluators are paid by the pool and tend to be permissive in what they consider essential. More patents in the pool means more revenue for the pool and its administrator. This doesn't mean the evaluation is corrupt, but the incentive structure is worth understanding.

Patent pools raise antitrust questions — competitors collectively setting a price for their patents looks like price-fixing. Courts and regulators have generally approved patent pools for standards when they include only necessary claims, when they are non-exclusive (patent holders can still license independently), and when they don't extend to non-essential patents. The key distinction is that pooling necessary claims for a standard reduces transaction costs without restricting competition, while pooling broader patent rights could suppress it.

For practitioners, the existence of a patent pool may be the first concrete answer to a client's question: "what will it cost to implement this standard?" If a pool exists, it provides a published rate. If it doesn't, the answer is the familiar: implement and be prepared for someone to knock on your door.

7.6 Over-Disclosure and Strategic Behavior

In RAND regimes, a significant practical problem is **over-disclosure** — patent holders disclosing far more patents than are actually essential.

Because disclosure obligations are based on a reasonable belief that a patent might be a necessary claim — not a formal determination — patent holders have every incentive to disclose broadly. If you're planning to license, disclosing early establishes your position. If you're not sure, it's safer to disclose than to risk the consequences of non-disclosure. Note that this is a RAND-specific issue. In royalty-free organizations, there is generally no need for patent disclosure since participants are getting a royalty-free license regardless — the relevant mechanism there is the exclusion process, not disclosure.

Over-disclosure creates real problems for the standards development process. A working group confronted with a large volume of patent disclosures may feel compelled to evaluate them, potentially triggering design-around efforts, slowing the development cycle, or diverting engineering resources from technical work to patent analysis. In some cases, over-disclosure is strategic — a way to signal the size of a patent position and influence the technical direction of the spec.

The result is that the disclosure list for a mature RAND standard can be enormous, with far more patents listed than are actually essential. Implementers looking at the disclosure list may overestimate their royalty exposure. Patent holders who later form pools will submit patents that were disclosed but may not survive independent evaluation.

For the standards body, managing the volume of disclosures is an ongoing challenge. For implementers trying to assess their exposure, over-disclosure creates noise that makes due diligence more difficult.

7.7 Standard-Essential Patents and the Cross-Licensing Reality

A term you'll encounter frequently in RAND discussions is **SEP** — standard-essential patent. A SEP is simply a patent that contains one or more claims that are necessary to implement a standard. It's the patent-level equivalent of the "necessary claims" concept discussed in Chapter 6. If a patent has at least one claim that is essential to a standard, the entire patent is often referred to as a SEP, even though not all of its claims may be standard-essential.

SEPs matter because they are the patents subject to RAND commitments. They are the patents that get disclosed, that get submitted to patent pools, and that are the subject of licensing negotiations and litigation. The term is used pervasively in the wireless, video codec, and telecom industries, and increasingly in policy discussions by regulators and legislatures.

How Licensing Actually Works

The textbook model of RAND licensing — patent holder offers a license, implementer negotiates terms, they reach agreement on a per-standard royalty — is an oversimplification of how licensing works among large companies.

In practice, SEP licensing between major companies is rarely done on a standard-by-standard basis. Instead, it happens as part of **broad cross-licensing agreements** that cover both standard-essential and non-essential patents across multiple technology areas. Company A licenses its entire portfolio to Company B, and Company B does the same in return. The agreement may cover SEPs for wireless standards, SEPs for video codecs, non-essential patents on implementation techniques, and patents entirely unrelated to standards — all in a single deal.

These cross-licenses are the dominant form of patent licensing among large technology companies. They are negotiated bilaterally, are almost always confidential, and often involve balancing payments where one party's portfolio is larger or more valuable than the other's. The RAND commitment provides the floor — neither party can refuse to license SEPs on reasonable terms — but the actual deal encompasses far more than SEPs alone.

Patent pools and standalone SEP licensing programs do exist and are significant in specific areas — particularly wireless (where pool-like structures like Avanci operate) and video codecs (where MPEG LA and Via Licensing administer pools). But for the largest companies, these are often supplementary to the broader cross-license relationships.

This distinction matters for practitioners because it means that the RAND commitment, while foundational, is one input into a much larger licensing dynamic. A client asking "what will it cost to license the SEPs for this standard?" may be asking the wrong question. The more relevant question is often: "what is the state of our overall patent relationship with the companies that hold SEPs in this space?"

7.8 RAND-RF Hybrids and Emerging Models

RAND-Z with RAND fallback is a model where participants initially commit to license at zero royalty, but retain the right to fall back to RAND (royalty-bearing) terms if they choose. In theory, this offers the best of both worlds — royalty-free as the default, with a safety valve for patent holders who need to monetize. In practice, these structures tend to produce a race to the bottom. The moment one participant elects the RAND fallback and begins charging, other patent holders follow — because there's no incentive to continue offering royalty-free terms when your competitors are collecting royalties on the same standard. The RAND-Z default effectively collapses into a RAND regime once anyone exercises the fallback.

Dual-mode organizations like JDF allow different working groups within the same organization to operate under different IPR modes. One working group might use royalty-free while another uses RAND, depending on the commercial dynamics of the technology and the participants involved. This flexibility avoids forcing a one-size-fits-all choice at the organizational level.

These hybrid approaches reflect a practical reality: the binary choice between RAND and royalty-free doesn't always map cleanly to the needs of a particular project. The trend is toward more flexibility in policy selection, allowing the IP terms to match the specific context of the work rather than the institutional default of the hosting organization.

7.9 The Practical Reality

For all the complexity of RAND policy terms, the practical reality for most implementers is straightforward: implement the standard and see what happens.

In the vast majority of cases, nobody comes knocking. Active licensing programs are concentrated in wireless, video codecs, and a handful of other areas. Outside those spaces, RAND commitments are largely theoretical — they exist in the policy, patents may be disclosed, but no one sets up a licensing program and no money changes hands.

When licensing does happen, it happens outside the standards body through bilateral negotiation. If the parties can't agree, the dispute goes to court — not back to the standards organization. The standards body's role ends with the commitment. Everything after that is between the patent holder, the implementer, and potentially a judge.

The most difficult conversation with a client in this space is the one where they ask: "How much will it cost to implement this standard?" In most cases, the honest answer is: we don't know. It could cost nothing. It could cost a meaningful amount. The uncertainty is inherent in the RAND model, and no amount of policy drafting eliminates it. What the policy does is ensure that if someone does come knocking, there's a framework — reasonable terms, non-discriminatory treatment — that constrains the negotiation. The alternative would be no framework at all, which would be worse.

Practice Tips

1. Disclosure is mandatory based on actual personal knowledge — not optional, regardless of what engineers may assume.
2. "Reasonable" rates are determined through negotiation or litigation, not by the standards body.
3. Ask about your client's overall patent relationships, not just per-standard SEP licensing — among large companies, licensing happens through broad cross-licenses.
4. Read the declaration form in addition to the policy — forms sometimes modify the underlying terms.
5. There is no duty to search, but failure to disclose a known patent creates significant enforcement problems later.

Chapter 8 — Royalty-Free Patent Policies and the Exclusion Mechanism

Royalty-free patent policies are the foundation of most modern technology standards work. The premise is straightforward: participants commit to licensing their necessary claims at no charge, creating a zone where anyone can implement the standard without worrying about patent royalties.

In practice, these policies are anything but straightforward. The details — scope, triggers, definitions, defensive provisions — determine whether the commitment is meaningful or hollow. This chapter walks through the anatomy of a royalty-free patent policy, using real structures to illustrate the choices and tradeoffs involved.

8.1 The Philosophy: Royalty-Free and Litigation-Free

A royalty-free patent policy creates two things simultaneously: a **royalty-free zone** and, at least among the participants, a **litigation-free zone**.

The royalty-free part is obvious — patent holders commit not to charge for licenses to their necessary claims. The litigation-free aspect is equally important, though its boundaries are worth noting. When everyone in the room has made a commitment to license their necessary claims at no cost, the incentive to assert those patents against fellow implementers drops to near zero. The result, among participants, is an environment where companies can implement the standard without the background threat of patent litigation from other members.

This protection doesn't extend to parties outside the tent. Non-participants who hold patents reading on the standard are not bound by the policy and remain free to assert. The royalty-free commitment creates safety among those who made it — not universal immunity.

This is the value proposition that makes royalty-free standards attractive, particularly in the software and web spaces where the technology moves fast and the cost of clearing patent rights for every implementation would be prohibitive.

One practical factor that reinforces adoption of royalty-free policies: familiarity. Organizations like JDF and OASIS offer multiple IPR policy modes, including the ability to adopt established policies like the W3C Patent Policy or the ISO/IEC patent policy by reference. The W3C policy is not a model of clarity — it can be genuinely confusing on a careful read. But it's one of the first royalty-free patent policies, and people know it. The value of a familiar policy is the same as the value of a familiar open source license. You could write a bespoke agreement that's technically superior to MIT or Apache 2.0. Nobody would trust it.

But if you say "we're using the W3C policy," the conversation moves to substance rather than stalling on the framework. In standards, where the goal is broad participation, practical trust often matters more than legal precision.

The tradeoff is that patent holders give up the ability to monetize their standard-essential patents — at least for that standard. Industries where companies have made significant R&D investments that they expect to recover through patent licensing tend to resist royalty-free policies. That's why telecom and video codec standards often use RAND rather than royalty-free. We covered that in Chapter 7.

8.2 The Licensing Commitment in RF Policies

As discussed in Chapter 3, a patent policy creates a commitment to license — not an actual license. The license itself is negotiated bilaterally outside the standards body. That general principle applies to royalty-free policies, but with two RF-specific nuances worth noting.

First, there are exceptions to the commitment model. A small number of organizations use actual licenses or non-assert agreements — the Open Web Foundation Agreement, for instance, is a direct non-assert commitment rather than a promise to negotiate. Some policies grant licenses that attach automatically upon implementation. But the commitment model remains the norm.

Second, in practice, the bilateral licensing step almost never happens in RF regimes. It is extremely rare for anyone to seek or offer a formal royalty-free license. When no money is changing hands, there's little incentive for either side to formalize the arrangement — and little practical remedy if they don't. If you've already committed to license at zero royalty, what are the damages? The commitment itself provides the assurance implementers need, and the bilateral license remains theoretical.

8.3 Scope: The First Thing to Get Right

The scope of a patent commitment is one of the most consequential decisions in setting up a standards project, and one of the most contentious to negotiate.

Organization-Wide vs. Working-Group-Specific

Patent policies apply at one of two levels: across the entire organization, or on a working-group-by-working-group basis.

Organization-wide policies are simpler to administer. Everyone who joins the organization makes the same commitment to all of its output. But they carry a significant risk: if the organization expands into new work areas — and organizations almost always expand — participants find themselves with patent commitments to specifications they weren't involved in developing, don't care about, and may not have anticipated.

The expression is that oxygen expands to fill the room. Give a standards organization the opportunity to expand, and it will. If your patent commitment covers everything the organization does, you're signing up for an unknown future.

Working-group-specific policies address this by confining the patent commitment to the output of the working groups you actually participate in. Each working group has its own charter, its own scope, and its own IP terms. You opt in by joining the working group. If a new working group forms to work on something unrelated to your interests, your patents aren't implicated.

This is why organizations like JDF and OASIS allow different working groups within the same organization to operate under different IPR modes — one group might be royalty-free while another, for strategic reasons, uses RAND. The flexibility to match the IP terms to the specific context of each work stream is a significant design advantage.

Getting Scope Right

Each working group should have a defined scope that sets the outer bounds of what the group can work on. This scope cascades into the patent commitment: if the group produces work outside its scope, the patent commitment may not cover it.

Scoping is more art than science, and it's one of the hardest things to get right.

Too narrow, and the group can't adapt as the work evolves. If the scope says "microwave ovens" and the group discovers halfway through that the real product should include a toaster function, they can't do it without rechartering — a process that can take months and require re-executing agreements.

Too broad, and participants face open-ended patent exposure. If the scope says "kitchen appliances that apply heat," a participant with valuable toaster patents may find those patents swept into a commitment they never intended to make.

Companies without significant patent portfolios in the relevant space tend to push for broad scopes — they have nothing to lose and want maximum flexibility. Companies with patents they intend to monetize or protect push for narrow scopes with explicit carve-outs. This tension is inherent in every scoping discussion, and it's one of the places where the negotiation gets real.

The contrast can be stark. When the Alliance for Open Media was formed to develop royalty-free video codecs, the scope of the video coding working group was essentially one line: we are going to create video codecs. That worked because the founding companies were aligned on the goal and had no intention of monetizing codec patents within the group. When the same organization later formed an audio codec working group, the scoping document was a page and a half of detailed carve-outs, because participants had different strategic positions and needed to protect specific aspects of their portfolios.

8.4 Who Makes the Commitment

Patent policies specify who is bound by the commitment. This has two dimensions: which entity within a corporate structure, and what triggers the obligation.

Affiliates and Corporate Structure

Most policies extend the patent commitment to patents "owned or controlled by a member or its affiliates." The intent is to prevent companies from hiding standard-essential patents in subsidiaries or affiliated entities to avoid the commitment.

This seems straightforward until you encounter real corporate structures. A venture-backed startup may be controlled by a VC that also controls hundreds of unrelated companies. A conglomerate's participation through one division doesn't necessarily mean it can — or should — commit the patents of its nuclear reactor division. An entity sitting several layers deep in a corporate hierarchy may have legitimate reasons why it can't bind its parent's entire portfolio.

These situations arise frequently and there's no universal answer. Sometimes you accept that the participant can only commit the patents they directly control. Sometimes you require them to sort out their internal structure before participating. Sometimes you negotiate a side letter that clarifies the scope of the commitment. The key is to understand the issue and address it explicitly rather than discovering a gap after the standard is published.

After-Acquired Patents

Most well-drafted policies specify that the commitment covers patents owned "now or at any time in the future." This means that if a participant acquires a patent after making the commitment — whether through R&D, acquisition, or any other means — that patent is pulled into the commitment if it reads on the standard.

Without this provision, a participant could make a commitment, then acquire a blocking patent the next day and assert it freely. The "now or any time in the future" language closes that gap. It seems obviously necessary, and it is, but it has been the subject of disputes.

8.5 Who Gets the License

A royalty-free commitment can run to different recipients, and the choice matters.

The most common approach is that the commitment extends to **anyone implementing the standard** — the entire world. This maximizes the value of the standard by ensuring universal access. It also means that non-participants — companies that didn't join the working group, didn't contribute, and didn't make any reciprocal commitment — still benefit from the patent commitment.

Some organizations restrict the commitment to **members only**. The rationale is that participants should have skin in the game: if you want the benefit of the patent commitment, you should join the organization, contribute to the work, and make your own commitment in return. This also serves as a membership incentive — joining the organization is the only way to get the royalty-free license.

The member-only approach has a practical downside: it limits adoption. If the goal is ubiquitous implementation of the standard, restricting the patent license to members creates a barrier. On the other hand, if the goal is to build a committed ecosystem of participants who all contribute and all benefit, the member-only model can make sense.

8.6 Reciprocity and Defensive Termination

Two provisions commonly appear in royalty-free policies that create conditions on the commitment: reciprocity and defensive termination.

Reciprocity

Reciprocity means that the royalty-free commitment is conditioned on the recipient also offering royalty-free terms back. If you want a license from me, you have to offer one to me as well.

This sounds fair, but the details get complicated. Does reciprocity mean "back to me" or "back to the world"? If I license my patents to you royalty-free, and the reciprocity provision says you must license yours back to me, that helps me but not the broader ecosystem. If reciprocity means you must make the same commitment to everyone, it effectively extends the patent commitment to non-participants who license from a participant — a much broader obligation.

How reciprocity is enforced in practice is even less clear. The bilateral license is negotiated privately, and the specific terms are generally confidential. A third party relying on a reciprocity chain has limited visibility into what was actually agreed.

Not all royalty-free policies include explicit reciprocity provisions. Some take the position that it's a perfectly reasonable bilateral negotiation term and doesn't need to be mandated by the policy. Others build it in as a standard feature.

Defensive Termination

Defensive termination allows a patent holder to revoke their royalty-free commitment against a party that asserts patent claims against them. The logic is straightforward: if you sue me for patent infringement related to this standard, I should be able to revoke my commitment to you.

This provision exists in most royalty-free policies and is generally uncontroversial in principle. The details — what triggers it, how broad the revocation is, whether it's limited to the asserting party or extends more broadly — vary across policies and can be significant.

Defensive termination serves an important deterrent function. It raises the cost of patent assertion within the standards ecosystem, because any assertion risks losing the royalty-free licenses the asserter depends on from other participants.

8.7 Comparing Major RF Policies

While the principles above are common across royalty-free policies, the details differ in important ways. Three frameworks dominate the landscape.

The JDF IPR Policy

JDF offers multiple IPR modes that a working group can select when chartering. The royalty-free mode provides a participation-triggered commitment — joining the working group creates the obligation, not the act of contributing. Scope is confined to the working group's charter. Exclusions follow a rolling model, placing the onus on the participant to track the spec and act before finalization. The policy is intentionally lightweight, designed for small to mid-size groups of like-minded participants where exclusions are expected to be rare.

The W3C Patent Policy

The W3C policy, first adopted in 2004 and updated several times since, was one of the first formal royalty-free patent policies. It uses a participation-triggered commitment tied to working group membership. A key feature is its formal exclusion process: when a specification reaches a defined maturity stage (Last Call or Candidate Recommendation), participants have a 150-day window to review and exclude. This is longer than most other policies and reflects W3C's emphasis on giving participants adequate time to review complex web specifications. The W3C policy applies on a working-group-by-working-group basis, and the commitment covers the specific Recommendation produced by that group. One area of ongoing debate is whether the commitment extends to subsequent versions of a specification — W3C has generally taken the position that new versions require new commitments.

The OASIS IPR Modes

OASIS provides three IPR modes: RAND, RF on RAND Terms, and RF on Limited Terms. Each OASIS Technical Committee selects its mode when chartered, and the mode governs all work produced by that committee. The RF on RAND Terms mode is the most commonly used for technology standards — it requires royalty-free licensing but otherwise follows RAND-style mechanics, including a formal disclosure and exclusion process. The RF on Limited Terms mode is more restrictive, limiting the scope of the commitment more narrowly. OASIS uses a formal call for patents tied to defined specification stages (Committee Specification, OASIS Standard), giving participants a structured window for exclusion. The specificity of OASIS's staged process makes it well-suited for organizations that want clear procedural guardrails.

Key Differences at a Glance

Feature	JDF	W3C	OASIS
Commitment trigger	Participation	Participation	Participation
Exclusion model	Rolling	Formal call (150-day window)	Formal call (tied to spec stages)
Scope	Working group charter	Working group	Technical Committee
Multiple modes	Yes (RF, RAND, W3C, ISO/IEC)	No (RF only)	Yes (RAND, RF/RAND, RF/Limited)
Versioning	Configurable	New commitment per version	Tied to specific spec version

8.8 Patent Exclusions: The Safety Valve

Any royalty-free policy that doesn't provide an exclusion mechanism should be a red flag. The ability to exclude specific patent claims from the royalty-free commitment is the safety valve that makes participation possible. Without it, the theory is that companies with significant patent portfolios simply wouldn't join.

Why Exclusions Exist

When you participate in a royalty-free working group, you're committing to license your necessary claims at no charge. That's a meaningful concession. The exclusion mechanism gives you a way to protect specific patents that you're unwilling to commit — either because the spec evolved in a direction you didn't anticipate, because you have a change in business circumstances, or because you have particular claims you need to preserve for other purposes.

In practice, exclusions are rare. Most participants in royalty-free organizations never exclude anything. The shared expectation — sometimes explicit, sometimes not — is that participants are entering a royalty-free zone and will behave accordingly. Exclusions are widely considered the nuclear option.

How Exclusions Work

Exclusion mechanisms come in two forms.

Call for patents. The organization issues a formal notice — typically when a specification reaches a near-final draft — declaring that participants have a defined window (usually 30 to 60 days, often settling at 45 after the usual negotiation) to review the spec and declare any exclusions. This is a specific trigger with a specific deadline.

Rolling exclusion. There is no specific call for patents. Participants have a standing obligation to declare exclusions before the specification is finalized. There's no formal trigger —

the onus is on the participant to track the spec's evolution and act in time.

Both approaches have tradeoffs. A call for patents gives participants a concrete deadline and a piece of paper that prompts internal review — particularly valuable in large organizations where the standards team may not know what the patent team holds. A rolling exclusion is lighter on the organization and works well when participants are engaged enough to track the work themselves. For smaller, like-minded groups where exclusions are expected to be rare, rolling exclusions reduce administrative burden. For larger or more adversarial environments, a formal call for patents provides better discipline.

What You Must Disclose When Excluding

When a participant makes an exclusion, they typically must identify the patents or patent applications being excluded and the portions of the spec they believe are implicated. The level of detail required varies across policies and is itself a negotiation point.

For issued patents, most policies require at minimum the patent number and the relevant spec sections. For unpublished patent applications, the requirement is necessarily different — the application isn't public, so the policy may ask for a description of the subject matter rather than specific numbers.

Some organizations are moving toward more rigorous exclusion requirements: specific claim identification, mapping to spec sections, grouping by patent family, and even a good-faith attestation that the excluded claims are indeed necessary claims. The motivation for these enhanced requirements is experience with participants who have abused the system by mass-excluding hundreds of patents without meaningful review.

When Exclusions Go Wrong

Exclusions done in good faith, narrowly focused on one or two specific claims, with engagement and communication with the working group, are manageable. The working group can evaluate whether to design around the excluded claims or accept them.

Mass exclusions — where a participant searches their patent database by keyword and dumps hundreds of patent numbers — are a different matter. Standards bodies are not equipped to process that volume. It's highly disruptive, it undermines the royalty-free commitment, and it's generally regarded as bad faith. But if the policy allows it, the organization has limited recourse.

This is an area where policies are actively evolving. Newer versions of IPR policies are adding procedural requirements specifically designed to raise the cost of bad-faith exclusions — requiring more detailed disclosures, good-faith attestations, and in some cases limiting the scope of what can be excluded against existing spec material that has already been through a prior exclusion window.

Versioning and Exclusions

One of the most contentious questions in exclusion practice is whether a new participant can exclude against material from prior versions of a specification.

If a working group produced version 1.0 under a royalty-free commitment, and a new participant joins to work on version 2.0, can they exclude claims that read on the version 1.0 material that carries forward into 2.0? Reasonable positions exist on both sides.

The argument for allowing it: the new participant only committed to the working group when they joined. They shouldn't be bound by past versions they had no hand in developing.

The argument against: the new participant joined with full knowledge of the existing spec. They had the opportunity to review it. Allowing exclusions against entrenched, widely-deployed technology is massively disruptive to an ecosystem that relied on the royalty-free commitment.

Policies handle this differently. Some are explicit that the commitment covers only the current version under development. Others use formulations like "substantially similar purposes" to extend coverage to related prior work. Many are silent, which creates ambiguity that surfaces at the worst possible time.

Common pitfalls that apply across all patent policy types — scope creep, versioning gaps, withdrawal mechanics, and more — are covered in Chapter 10.

Practice Tips

1. Read the exclusion mechanism closely — distinguish between formal calls for patents and rolling exclusions.
2. When excluding, identify specifically which patents and which spec sections are implicated.
3. Check whether the commitment extends to anyone implementing the standard or to members only.
4. Verify how the policy handles new participants joining mid-process and prior-version material.
5. Map scope carefully — if a new version is treated as a new commitment, older versions may lose patent protection while still being implemented.

Chapter 9 — Non-Asserts, Pledges, and Covenants

The patent policies discussed in Chapters 7 through 8 operate within standards bodies — they're part of the governance framework that participants agree to when they join. Non-asserts and pledges started as something different: unilateral commitments made outside the standards process, typically by a single company, promising not to assert certain patents against implementers of a specification.

Over time, however, non-assert agreements have been adopted by standards bodies themselves. Organizations like OASIS and JDF now include non-assert options alongside their traditional RAND and royalty-free IPR modes. The Open Web Foundation agreements, originally developed as standalone instruments, are used by some organizations as their primary IPR framework. What began as an alternative to the standards process has become part of it.

Understanding the history of these agreements is useful for understanding why they exist, what problems they solve, and what their limits are.

The core motivation was distrust of RAND-based commitments. Under a RAND regime, the commitment is to offer a license on reasonable terms — but you don't see the actual terms until you go negotiate. For implementers, particularly in the open source community, this created a "gotcha" risk: you build to the standard, then discover that the license terms are unacceptable — incompatible with your distribution model, laden with conditions, or simply too expensive. By that point, you're already committed to the technology.

Non-asserts addressed this by putting the terms up front, before implementation. The rights are specified in the agreement itself — not deferred to a future bilateral negotiation. This was a model open source developers were already accustomed to. You read the license, you know what you're getting, you implement accordingly. No surprises after the fact.

9.1 Beyond Traditional Patent Licensing

A patent license grants specific rights to use patented technology. A non-assert is a promise not to sue. For a long time, the legal community debated whether these were functionally different — does it matter if I grant you a license to drive my car versus promise not to sue you when you drive it?

The Federal Circuit eventually concluded that for patent purposes, the distinction is largely academic. Both achieve the same result: the implementer can practice the technology without fear of enforcement. The original motivation for the non-assert formulation was to avoid patent exhaustion — the doctrine that once a patent holder licenses a product, their rights are exhausted for downstream recipients. By calling it a "promise" or "covenant" rather than a "license," drafters hoped to preserve the ability to assert against downstream parties. Courts were not impressed by the distinction.

Despite this, the non-assert formulation persists. Some of these instruments are drafted as promises, some as covenants, some as pledges. The terminology varies, but the practical effect is the same: the patent holder is committing not to enforce specified patents against specified implementations.

9.2 A History of Non-Assert Agreements

The evolution of non-assert agreements tracks the broader history of patent politics in the technology industry. Each generation was a response to specific trust failures and competitive dynamics.

The Sender ID Episode

In 2004, an early and instructive episode involved Microsoft's Sender ID email authentication specification at IETF. Microsoft made a royalty-free patent declaration — which sounded generous. But the actual license terms, linked from the declaration rather than included in it, imposed conditions that were incompatible with open source distribution models. The license was explicitly non-transferable and non-sublicensable, required formal execution, and distinguished between categories of users in ways that conflicted with how open source software is distributed. To make matters worse, the underlying patent application wasn't publicly available, so participants couldn't even evaluate the scope of the claims they were being asked to accept.

Major open source projects, including the Apache Software Foundation and the Debian project, publicly refused to implement the specification. The IETF working group itself ultimately rejected the proposal, concluding that the patent encumbrances could not be ignored and that participants could not accurately assess patent claims that weren't public.

The episode received mainstream press coverage and generated lasting distrust. It became a reference point every time Microsoft tried to engage with the open source and standards communities in the years that followed. But it also served as an important internal lesson. The Sender ID experience was one of several milestones that eventually led Microsoft to fundamentally rethink its relationship with open source — a process that took years but ultimately resulted in the company becoming one of the largest contributors to open source and a proponent of royalty-free standards. For other companies navigating similar tensions, it's a useful reminder: the terms you attach to a royalty-free commitment will be scrutinized, and getting them wrong can set back your credibility for a decade.

The Document Format Wars: ODF and OOXML

The mid-2000s document format standardization battle produced the most consequential wave of non-assert agreements.

Sun Microsystems, which had acquired StarOffice and was championing the Open Document Format (ODF), issued a non-assert for its patents reading on the ODF specifications developed through OASIS. The commitment was irrevocable, tied to specific versions where Sun participated, and included a defensive suspension provision. It was straightforward and well-received by the open source community — partly on its merits, partly because Sun had significant credibility with that community at the time.

IBM took a different approach. IBM published a "Statement of Non-Assertion of Named Patents Against OSS" — a pledge covering 500 named patents. The open source community praised it. On closer reading, however, many of the 500 patents were of marginal value, and the defensive suspension clause was exceptionally broad: if you asserted any patent against any open source software — regardless of whether it involved IBM's patents or IBM at all — the pledge was revoked. The actual legally binding terms were buried at the end of a very long document listing patent numbers. It was a significant public relations achievement with modest practical substance.

The Open Specification Promise

In response to criticism that its royalty-free commitments couldn't be trusted, Microsoft developed the Open Specification Promise (OSP). The OSP was deliberately drafted as a "promise" rather than a license, in part to avoid patent exhaustion and in part to avoid GPL compatibility concerns. It was structured as a direct commitment from Microsoft to each individual implementer — not a chain of rights that passed from one party to the next.

The OSP allowed subsetting — implementers could implement partial specifications and still receive patent coverage. This made sense for the office file formats, where nearly everything was optional, but created concerns when applied more broadly: a subset of a spec taken out of context could be used for purposes the original standard never intended.

Over time, the OSP became a default agreement applied to many specifications beyond its original purpose. As the number of specs covered grew, the defensive suspension provision — which revoked the promise if you sued over any covered specification — became increasingly broad.

The Open Web Foundation Agreements

The Open Web Foundation (OWF) agreements emerged from the open source community's dissatisfaction with all of the above. The goal was to create a community-developed, non-assert-style agreement specifically designed for standards and specification licensing that would be unambiguously compatible with open source.

The OWF agreements have two components: a Contributor License Agreement (CLA) for participants contributing to the spec, and a Final Specification Agreement (FSA) that grants patent and copyright rights to implementers of the final specification. The structure mirrors open source licensing patterns — familiar territory for the communities that use them.

The OWF agreements have been adopted by numerous specifications and, importantly, by standards organizations themselves — including as an IPR mode option within organizations like OASIS and JDF. This institutional adoption is what moved non-asserts from being ad hoc unilateral commitments to being a standard part of the standards toolkit. Their value is largely the same as any well-known, community-developed agreement: trust. They aren't proprietary to any single company, they were developed through an open process, and they're understood by the communities that rely on them.

9.3 The Community Specification License

The **Community Specification License (CSL)** represents the next evolution — taking the principles behind the OWF agreements and packaging them for the way specifications are actually developed today: in Git repositories, with asynchronous collaboration, using lightweight governance.

Origins and Design Philosophy

The CSL was developed by the Joint Development Foundation and is maintained under the Linux Foundation. It draws on the OWF agreements and the Alliance for Open Media patent agreements, distilling their IP terms into a form that can be adopted simply by forking a template repository.

The design philosophy is that starting a collaborative specification project should be as frictionless as starting an open source project. Fork the repo, define your scope and governance, and begin working. The IP terms — copyright licenses, patent grants, exclusion mechanisms — are already built into the framework.

How CSL Differs from Traditional IPR Policies

Unlike the patent policies discussed in Chapters 7 and 8, the CSL is not embedded in a standards organization's governance. It's a standalone set of agreements that any group can adopt, with or without a formal organizational structure.

Contributors grant a perpetual, royalty-free, non-exclusive copyright license and a patent license covering necessary claims. Patent exclusions are permitted but must be declared. The commitment is contribution-triggered — you make the commitment by submitting material to the repository.

The CSL is designed for specifications, not source code. This is an important distinction. Open source licenses grant rights to specific codebases. The CSL grants rights for independent implementations of a specification — a broader scope that ensures any implementer, using any codebase, receives patent coverage for practicing the spec.

When to Use CSL vs. a Full Standards Framework

The CSL is the right tool when you need IP clarity for collaborative spec development but don't need the full machinery of a standards organization — no corporate entity, no membership tiers, no formal governance structure. It works well for small collaborations, early-stage specification work, and projects that may later move into a formal standards body.

It also provides a natural on-ramp. A project that starts under the CSL can later transition to JDF or another standards framework if it outgrows the lightweight model. The IP terms are compatible, and the transition doesn't require re-doing the patent commitments from scratch.

The CSL is not the right tool when you need formal governance, tiered membership, certification programs, or the institutional weight of a recognized standards body. For those, the full frameworks discussed in Chapter 2 are more appropriate.

9.4 Anatomy of an Effective Non-Assert

Whether you're evaluating an existing non-assert or drafting one, several elements determine whether the agreement is meaningful or merely theatrical.

Scope

What specifications does the commitment cover? A specific version, or all versions in perpetuity? Only the mandatory portions, or optional portions as well? Only conformance implementations, or subsets? Each of these choices has consequences, and the answers should match the goals of the commitment.

Irrevocability

Is the commitment irrevocable? Most modern non-asserts are, but older ones sometimes include conditions or expiration provisions that limit their durability. An irrevocable commitment provides the strongest assurance to implementers.

Sublicensing and Exhaustion

Can the rights be passed downstream? This matters for supply chain scenarios where components implementing the spec are incorporated into larger products. If the commitment is structured to avoid exhaustion — as some early agreements attempted — downstream recipients may not be covered.

Defensive Termination

Under what circumstances can the commitment be revoked? A narrow defensive termination — revoking the commitment only against a party that asserts patents on the same specification — is reasonable and widely accepted. A broad defensive termination — revoking against anyone who asserts any patent against any open source software — is a different proposition entirely, and the implications grow with every specification added to the agreement.

Why Your Own Agreement May Be Technically Better but Practically Worse

This is one of the most important practical lessons in standards IP. You can draft a non-assert or patent pledge that is, on every technical dimension, superior to the OWF agreements or the CSL. Better definitions. Cleaner scope. More precise defensive termination.

Nobody will trust it.

The value of a well-known, community-developed agreement is that people recognize it. They can go to their legal team and say: this is the OWF agreement, it's been used by dozens of specs, here's the analysis we did last time. The conversation is about substance, not about whether the framework itself is trustworthy.

A bespoke agreement, no matter how well drafted, starts from zero trust. Every potential participant's legal team has to review it from scratch. They'll look for traps. They'll compare it unfavorably to what they know. The friction this creates is real, and it's often the difference between a specification that gets broad adoption and one that doesn't.

Use established agreements when you can. Save bespoke drafting for situations where the established options genuinely don't fit.

Practice Tips

1. Evaluate every non-assert on four dimensions: scope, irrevocability, sublicensing/exhaustion, and defensive termination breadth.
2. Use established community-developed agreements (OWF, CSL) over bespoke drafts — trust matters more than technical superiority.
3. Watch defensive termination scope — it grows more problematic with every specification added to the agreement.
4. For lightweight projects, the Community Specification License provides IP terms without organizational overhead.

Chapter 10 — Patent Policy in Practice: Pitfalls, Applications, and Theater

The preceding chapters covered patent policy frameworks — necessary claims, RAND, royalty-free, exclusions, non-asserts. This chapter is about what happens when those frameworks meet reality. The pitfalls that trip up practitioners, the real-world examples that illustrate how these tools get deployed, and the role of theater in making patent commitments land.

10.1 Common Pitfalls Across All Patent Policies

These issues apply regardless of whether the policy is RAND, royalty-free, or non-assert. They don't surface in a casual read of the policy but can create real problems when they arise.

Scope Creep Without Rechartering

A working group starts with a narrow scope, the work expands organically, and nobody goes back to update the charter. The result may be specifications that fall outside the original scope — and therefore outside the patent commitment. This creates a spec without reliable patent coverage.

The risk is the same in both RF and RAND regimes. In RF, it means implementers may not have the royalty-free protection they assumed. In RAND, it means the disclosure and licensing framework may not apply to portions of the spec that grew beyond the charter.

The fix is straightforward: recharter when the scope changes. The reason it doesn't happen is that rechartering triggers a new review of patent commitments, and nobody wants to slow the technical work down. So the scope drifts, and the gap between the charter and the actual work grows until someone notices — usually at the worst possible time.

Participant vs. Member Ambiguity

Who exactly has made a commitment? Many policies distinguish between different tiers of participation — steering committee members, general members, observers, invited experts — and the patent commitment may apply differently to each. Some tiers carry full commitments. Others carry none.

Make sure you know which tier your client is in and what that means for both their outbound commitments (what they're giving) and their inbound rights (what they're receiving). A participant who joins as an observer to evaluate the work may discover they've made a patent commitment by virtue of attending, or conversely may discover they have no patent protection for their implementation because observers aren't covered by the commitment pool.

Contribution-Triggered vs. Participation-Triggered Commitments

Some policies require a patent commitment from anyone who participates in the working group, regardless of whether they contribute technology. Others only require a commitment tied to your actual contributions.

The difference is significant. A participation trigger means you're making a commitment by joining. A contribution trigger means you can observe without committing — but the commitment on what you do contribute may be broader, sometimes extending to anything in the final spec that relates to your contribution, not just the contribution itself.

Understanding which trigger your policy uses is essential for advising clients on what joining a working group actually means.

Versioning Gaps

The commitment covers the version of the specification that was developed under the policy. But what about the next version? Some policies carry the commitment forward to subsequent versions of the same spec. Others treat each version as a new commitment. If the policy is silent on versioning, you may have a commitment to version 1.0 and no commitment to version 1.1 — even if the changes are minor.

This applies to both RF and RAND commitments. In long-lived specifications that evolve through many versions, the versioning question can determine whether the patent protection that implementers rely on actually covers the version they're implementing.

Withdrawal Mechanics and Zombie Commitments

Most policies allow participants to withdraw from a working group, but the commitment to specifications completed before withdrawal survives. This is the "zombie commitment" — the participant is gone, but their patent obligations persist for the work that was finalized while they were there.

Understanding what obligations persist after withdrawal — and what the withdrawal process actually requires — is essential before your client joins. Some policies have specific notice periods. Others require a formal declaration. Some are silent on the mechanics, which creates ambiguity about when the withdrawal is effective and what commitments survive.

Affiliate and Successor Ambiguity

The question of which entities within a corporate structure are bound by the commitment is often under-specified. As discussed in Chapter 6, "owned or controlled" language is meant to prevent patent hiding, but the edge cases — VC-backed companies, conglomerates, entities acquired after the commitment — can be genuinely difficult.

This is equally critical in RAND contexts, particularly when companies undergo M&A transactions. If your client acquires a company that holds patents reading on a standard they've committed to, do those patents fall under the commitment? The answer depends on the policy language — and policies vary.

Normative References to Other Specs

A specification often doesn't exist in isolation. It normatively references other specifications — effectively saying "for this part, go implement that other spec and come back." When you follow a normative reference, you're subject to the referenced spec's patent policy, not your own working group's policy.

This can create unexpected results. A royalty-free specification that normatively references a RAND specification has effectively introduced a royalty-bearing element. An implementer who assumed everything was royalty-free discovers that one component of the standard carries licensing obligations under a different policy.

The fix is to be intentional about normative references at the drafting stage. Understand the IP terms of every referenced spec. If the policies aren't compatible, either find an alternative reference or make the dependency informative rather than normative.

The Form-Versus-Policy Trap

In some organizations, particularly international ones, the patent policy is one document and the declaration or disclosure form is another. The form sometimes includes terms that go beyond — or subtly modify — what the policy says. Additional reciprocity conditions, assignment obligations, or definitional modifications can appear in the form without being reflected in the policy text.

This is a trap for practitioners who read the policy and stop there. Always read the form. If the form modifies the policy, understand whether those modifications are binding and how they interact with the policy terms your client is relying on.

Mixed-Mode Organizations

Organizations that allow different working groups to operate under different IPR modes create a particular challenge: what happens when work from an RF working group and work from a RAND working group need to be combined? The patent commitments are different. The participant pools may be different. The terms under which the combined spec can be implemented may be unclear.

This doesn't arise often, but when it does, it can be genuinely difficult to resolve. The best mitigation is clear charter language that anticipates cross-group dependencies and specifies how the IP terms interact.

Late Joiners and Accumulated Obligations

A participant who joins a working group midway through the process faces a different situation than one who was there from the beginning. The spec is partially developed. Decisions have been made. The scope may have already drifted from the charter. And the patent commitment typically covers the final deliverable — including all the work done before the new participant arrived.

Some policies give late joiners a review period — typically 30 to 60 days — to evaluate the existing work and decide whether to stay or withdraw before committing. Others don't, meaning the commitment attaches immediately upon joining.

For practitioners advising a client who's considering joining an ongoing project, the key questions are: what's already been developed, does the current work fall within the chartered scope, and does the policy provide any on-ramp for new participants to evaluate their exposure before committing?

10.2 Real-World Applications

Theory only gets you so far. Several recent examples illustrate how patent commitments work — and don't work — in practice.

Google's Agent-to-Agent Protocol

Consider a unilateral spec release that, from a legal perspective, provided minimal patent coverage — but from a theater perspective, hit all the right notes.

The specification was released as an "open protocol" with contributions from dozens of technology partners. The language was deliberately chosen: open, collaborative, community-driven. The spec was published under an Apache 2.0 license with a standard Apache Contributor License Agreement (CLA) for contributions.

From a standards patent perspective, this provides almost no coverage. An open source license covers contributions — the specific code or text that individual contributors submitted. It does not cover the specification as a whole. If someone holds a patent on the interoperability mechanism described by the spec, the Apache license doesn't help. The patent grant in an Apache

CLA is tied to the contributor's contributions, not to independent implementations of the spec. This is fundamentally different from a standards patent commitment, which covers necessary claims across the entire specification regardless of who contributed what.

Yet the release was broadly praised. Partners provided enthusiastic endorsements. The press covered it favorably. Nobody raised patent concerns.

The takeaway is uncomfortable but important: the market doesn't always distinguish between legally robust patent coverage and the appearance of openness. A spec released with the word "open" used liberally, a recognizable license, and strong partner endorsements can achieve broader adoption than a spec with technically superior patent terms but less effective messaging. Getting the legal framework right still matters — but it's not sufficient on its own.

The Lesson

For practitioners, the lesson is that you need to operate on both dimensions simultaneously. Get the legal framework right, because it provides the foundation for everything else. But also understand how the framework will be perceived, communicated, and received. Ask yourself: if a competitor wanted to attack this, what would they say? If a journalist wrote about it, what would the headline be? If a skeptical engineer read the terms, what would they flag?

The best outcome is a legal framework that's both technically sound and publicly credible — which usually means using recognized, community-developed agreements rather than proprietary ones, and communicating the terms clearly rather than burying them in footnotes.

The Open Source License Misapplication Problem

The A2A example highlights a broader pattern worth noting: the increasing use of open source licenses for specification releases. Apache 2.0, MIT, and Creative Commons are familiar, trusted, and easy to apply. But they were designed for code and content, not for specifications.

An open source license grants rights to the specific work it covers — the code, the text, the file. A standards patent commitment grants rights to implement the described technology regardless of the specific expression. These are fundamentally different grants. Applying an open source license to a spec gives you copyright coverage (the right to copy and distribute the document) and perhaps a narrow patent grant tied to the specific contribution, but it does not give you the broad patent commitment that a standards policy provides.

This distinction matters most when the spec describes something that's independently implementable — a protocol, an interface, an interchange format. Anyone can read the spec and build their own implementation from scratch. If the only patent grant is tied to the contributor's specific expression, the independent implementer has no patent protection at all.

The trend toward using open source licenses for specs isn't wrong in all cases — for simple specs with no patent concerns, it may be perfectly adequate. But for specs that describe patentable technology, practitioners should ensure that the patent coverage matches the need. Open source familiarity is not a substitute for standards-appropriate IP terms.

10.3 The Theater of Standards: Signaling, Perception, and Trust

A theme running through this entire book is that standards work involves a significant element of theater.

This is not a pejorative observation. It's a description of how multi-party collaborations actually function. When dozens of companies are making patent commitments, evaluating IP risk, and deciding whether to implement, the decision isn't purely legal. It's about trust, signaling, and perception.

Signaling

The choice of patent policy, the choice of standards body, and the choice of licensing framework all send signals. Using a well-known, community-developed agreement signals trustworthiness. Using a proprietary agreement signals control — which may be appropriate in some contexts but will generate skepticism in others. Choosing a royalty-free framework signals commitment to broad adoption. Choosing RAND signals that patent monetization is on the table.

These signals matter because most participants don't read the actual policy terms in detail. They rely on heuristics: is this a framework I recognize? Is this hosted by an organization I trust? Has this been used before without problems? The signals answer those questions faster than a legal review ever could.

Perception Management

When things go wrong — a patent surfaces, an exclusion is filed, a participant withdraws — the response needs to address perception, not just legal substance. Pointing to the existing policy and saying "the framework already handles this" may be legally correct and practically insufficient. If the community doesn't trust the framework, or doesn't trust the party invoking it, the legal answer doesn't resolve the problem.

I've seen this play out firsthand. In one case, an existing royalty-free policy already covered the work. Legally, there was no gap. But the trust deficit required a response that went beyond legal adequacy — a public commitment under a recognized, non-proprietary agreement, endorsed by senior leadership, published where the community could see it. The law was fine. The perception wasn't.

Why Trust Matters More Than Perfection

You can draft a patent commitment that is, on every technical dimension, legally superior to the alternatives. If nobody trusts the drafter, it won't matter.

This is the same lesson that applies to non-assert agreements (Chapter 9), to the choice of IPR policy mode (Chapter 8), and to organizational design (Chapter 2). Trust is the currency of multi-party collaboration. It's built slowly through consistent behavior, and it's destroyed quickly through surprises — even unintentional ones.

For practitioners, the implication is that legal analysis is necessary but not sufficient. You need to understand how your advice will land in the community. You need to anticipate how the choice of framework, the choice of language, and the choice of public positioning will be received by people who have seen promises broken before. And sometimes, the best legal answer is not the best answer.

The Launch Announcement

One specific application of theater that practitioners should understand: the launch announcement.

When a new standard or specification is released, the announcement is often as important as the spec itself. It establishes the narrative. It identifies the partners. It frames the technology. And it signals the IP posture — whether this is "open," "collaborative," "royalty-free," or something else.

The launch announcement is where the legal terms get translated into language that the broader community can understand and react to. If the announcement says "open standard" but the IP terms are RAND, someone will notice the gap and call it out. If the announcement says "developed collaboratively with 60 partners" but the governance gives one company unilateral control, that tension will surface.

For practitioners advising on a spec launch, review the announcement with the same care you'd review the legal terms. Make sure the public positioning is consistent with the actual IP framework. Inconsistencies between the marketing and the legal terms are one of the fastest ways to generate the kind of trust-destroying scrutiny you're trying to avoid.

When Getting It Right Doesn't Matter

Perhaps the most uncomfortable aspect of theater in standards is the reality that sometimes, getting the legal framework right simply doesn't matter — and getting it wrong doesn't either.

Most standards-related patent commitments will never be tested in court. Most participants will never assert their committed patents. Most implementers will never seek or receive a formal license. The entire apparatus of patent policies, exclusion mechanisms, and licensing commitments exists as a framework for a negotiation that, in most cases, never happens.

This doesn't mean the framework is useless. It provides assurance. It creates expectations. It deters bad behavior. And in the cases where it does matter — the wireless licensing disputes, the video codec royalties, the occasional assertion by a company that changes its business model — the framework is essential.

But for the majority of standards engagements, the practical outcome is determined by trust, relationships, and community norms rather than by the specific terms of the patent policy. The policy sets the floor. The community determines the actual behavior. Understanding both — and advising your client on both

— is what separates competent standards counsel from someone who just reads agreements.

Practice Tips

1. Watch for scope creep without rechartering — out-of-scope work may lack patent coverage.
2. Read the declaration form in addition to the policy text — forms can subtly modify terms.
3. Check for versioning gaps — does the v1.0 commitment carry to v1.1?
4. Ensure launch announcements are consistent with actual IP terms — gaps will be noticed.
5. Don't assume an open source license provides the same patent coverage as a standards patent policy — they cover different things.

Chapter 11 — Standards Governance Structures

Governance in standards is where the theoretical frameworks of patent policies, organizational structures, and decision-making rules meet the practical reality of getting a group of competitors to agree on something. The governance structure determines who has power, how decisions get made, what happens when people disagree, and — critically — how the organization evolves over time.

This chapter covers the anatomy of standards governance, from boards and working groups down to the mechanics of chartering, membership design, and operational management.

11.1 Organizational Anatomy

Most standards organizations share a common structural pattern, though the terminology varies.

At the top sits a **governing body** — a board of directors, a steering committee, or in simpler structures, just the founding members acting collectively. This body sets the strategic direction, approves budgets, admits new members, and oversees the organization's operations. In contractual consortia, this function may be performed informally by the signatories to the agreement.

Below the governing body sit **technical committees** or **working groups** where the actual specification work happens. Each working group typically has its own charter defining its scope, its IPR mode, its deliverables, and its decision-making rules. The working group is where engineers draft text, debate technical choices, and iterate toward a final specification.

Some organizations add an intermediate layer — a **technical steering committee** that coordinates across multiple working groups, resolves cross-cutting technical issues, and ensures consistency between related specifications.

The relationship between these layers matters for IP purposes. As discussed in Chapter 8, patent commitments may apply at the organization level or the working group level. Understanding which layer your client's commitment attaches to is one of the first things to check when evaluating a new engagement.

11.2 Membership Models

How membership is structured determines who participates, who has influence, and who makes patent commitments. The design choices here have direct consequences for both governance and IP.

Open vs. Restricted Membership

Some organizations have fully open membership — anyone can join, anyone can participate, and everyone gets a voice. IETF is the extreme version: no formal membership at all, individual participation, and decisions by rough consensus. Most W3C working groups are open to any W3C member.

Others restrict membership, either by requiring an invitation, setting qualification criteria, or limiting the number of participants. This is common when a small group of companies wants to collaborate without opening the process to competitors they'd rather not include. The restriction has to be done carefully — antitrust law requires that similarly situated parties be treated alike, and exclusion needs objective justification.

Tiered Membership

Larger organizations typically use tiered membership. A common pattern:

Steering or founding tier. These members have the most influence — they sit on the governing body, approve specifications, and control the organization's direction. The top tier usually carries the highest dues and the greatest patent commitment obligations.

General or contributing tier. These members participate in technical work and have a voice in working groups, but may not have a vote on organizational governance. Their dues are typically lower.

Observer or associate tier. These members can follow the work and may attend meetings, but have limited or no voting rights. Some organizations use this tier as an on-ramp for companies evaluating whether to commit more deeply.

The naming conventions are inconsistent across the industry. "Founders" doesn't necessarily mean you were there at the founding — it's often just a label for the top membership tier. "Participant" may or may not carry patent commitments depending on the policy. Don't assume the label tells you the legal consequences. Read the agreement.

Permanent Board Seats and Succession

In smaller consortia, it's not uncommon for founding members to hold permanent seats on the board — seats that are guaranteed to the company for as long as they remain members, regardless of elections or rotations. This provides the founders with stability and control over the organization's direction, and it can be a powerful incentive for early commitment.

The interesting governance question is what happens when a founding member leaves. In some organizations, the seat reverts to the founding tier and is filled by another qualifying member. In others, the permanent seat simply disappears, and the board shrinks. In a well-designed structure, the vacated seat becomes an open seat available to any member at the appropriate tier — which can change the power dynamics significantly if the remaining founders now share the board with members who joined later under different expectations.

This is worth thinking about at the design stage. Permanent seats create stability in the short term but rigidity in the long term. If the organization's membership evolves — and it will — the governance structure needs to accommodate the change without requiring a constitutional crisis.

Who Is a "Participant" — Patent Commitment Implications

This is the practical punchline of membership design. The patent policy defines who makes a commitment and at what level. In some organizations, all members at every tier make the same commitment. In others, only the top tier commits, and lower tiers receive the benefit without contributing to the patent pool.

When advising a client on which tier to join, the patent commitment is often the decisive factor. A client with a significant patent portfolio may prefer a lower tier that doesn't trigger a broad commitment. A client with few patents may prefer the top tier for the governance influence, since the patent commitment costs them little. Understanding this dynamic — and being explicit with your client about what each tier means for their IP — is essential.

Case Study: DVD-CCA — Governance Complexity in Practice

The DVD Copy Control Association (DVD-CCA) is worth examining as an example of how complex governance structures can become when the stakeholders span fundamentally different industries.

DVD-CCA was established to manage and license the Content Scramble System (CSS) — the copy protection technology used on DVDs. It brought together three distinct industry groups: the entertainment industry (motion picture studios and content providers), consumer electronics manufacturers (companies building DVD players and drives), and the computer industry (hardware and software companies enabling DVD playback on PCs).

Each of these groups had different interests. The entertainment industry wanted strong copy protection. Consumer electronics manufacturers wanted interoperability and low licensing costs. The computer industry wanted flexibility to implement DVD playback in software.

To manage these competing interests, DVD-CCA used a governance structure with a twelve-member board — but the seats were not equally distributed. The entertainment industry held six seats, while consumer electronics and the computer industry held three each. This reflected the content owners' view that copy protection was fundamentally about their content, and that they should have the strongest voice in how it was managed. The imbalance created its own dynamics: the technology industries had to coordinate with each other to have any chance of outweighing the entertainment bloc on contested issues.

Separate membership classes for each sector were defined by objective qualification criteria — the entertainment industry class, for instance, required demonstrated box office revenue above a threshold, ensuring that the governance wasn't diluted by entities without meaningful stake in the outcome.

The voting structure added another layer of complexity. Rather than simple majority voting across all members, DVD-CCA used an **industry-by-industry voting model** where each of the three industry groups voted separately on certain matters. A proposal needed support within each group to advance — meaning the entertainment industry couldn't simply use its six-seat advantage to overrule the technology industries on every question, and a single industry group could effectively block progress on issues requiring cross-industry consensus.

Within the industry groups, voting reportedly followed a **sequential structure** — companies voted in a defined order, and the results were visible to other members. This created strategic dynamics that went well beyond the merits of any given proposal. If the first company in the order voted against a measure, it sent a signal to the rest of the group. Companies voting later could see which way the wind was blowing before committing. The order in which you voted — which might be determined by something as arbitrary as alphabetical order within the industry segment — could determine whether you were the one setting the tone or the one reacting to it.

This voting structure turned every decision into a multi-dimensional strategic exercise. Companies had to think not just about whether they supported a proposal on the merits, but about what their vote would signal to their industry peers, how the other industry groups would react, and whether a blocking vote from one group would create leverage for negotiations on other issues. Adding to the complexity, the order in which industry groups voted first rotated — so for each vote, a different industry was in the position of setting the initial tone. This rotation prevented any single industry from permanently controlling the signaling dynamic, but it also meant that strategic calculations shifted with every ballot. Alliances formed across industry boundaries. Companies that were competitors on the board found themselves voting together against a different industry group on copy protection scope, then opposing each other on licensing terms.

The structure also separated specification development from licensing and certification. The technical standards for CSS were developed through one process, while the licensing terms and compliance requirements were managed through another — each with its own governance dynamics.

This level of complexity was appropriate for DVD-CCA's specific situation: a technology that sat at the intersection of three powerful industries with conflicting incentives, significant money at stake, and content protection requirements that created genuine adversarial dynamics between content owners and technology implementers. It is not the norm for most standards engagements. Most projects don't need twelve-member boards with weighted industry-segment representation.

But DVD-CCA illustrates several governance principles that apply broadly:

1. When your stakeholders have genuinely different interests, the governance structure needs to account for those differences — through membership classes, board composition, or voting rules.

2. Unequal board representation creates its own dynamics — the minority groups must coordinate to be effective, which can either foster productive alliances or entrench adversarial positions.
3. The design of the voting structure shapes the strategic behavior of the participants — sequential visible voting creates different incentives than simultaneous secret balloting.
4. Separation of functions (spec development, licensing, certification) can reduce conflicts of interest and provide antitrust protection.
5. Objective membership criteria matter — without them, every admission decision becomes a political fight.
6. Complexity has a cost. The more elaborate the governance, the harder it is to operate, to amend, and to explain to new participants.

Dues and Financial Sustainability

One aspect of membership design that deserves attention is dues structure. Membership fees fund the organization's operations — staff, infrastructure, events, legal costs. The structure of these fees sends signals about who the organization is for.

Setting dues at the outset is more art than science. You don't know how much the organization will spend — that depends on how the work evolves, whether you hire staff, whether you hold events. You don't know how many members you'll attract — that depends on the technology's importance, the competitive dynamics, and whether the value proposition is compelling. So you pick a number, and you adjust later.

In practice, dues serve a function beyond revenue: they help companies **self-select** into the right tier of engagement. A high dues number at the steering level signals that this tier is for companies that are serious about the work and willing to invest. Companies that aren't prepared to commit at that level will self-select into a lower tier — or not join at all. This filtering function is often as important as the revenue itself. It ensures that the decision-making table is populated by participants with genuine commitment to the outcome.

The risk of setting dues too high is excluding participants whose technical contributions are valuable even if their budgets are limited — startups, academic institutions, smaller companies. Many organizations address this with sliding-scale dues based on company size or revenue, or with fee waivers for academic and nonprofit participants. Getting the balance right — enough revenue to sustain operations, enough filtering to ensure committed participants, low enough barriers to attract broad technical input — is one of the ongoing challenges of organizational design.

11.3 Forming a New Organization or Project

From the White Binder to Modular Frameworks

The history of standards formation is a story of progressive simplification.

In the early days, each new collaboration required a bespoke agreement negotiated from scratch. A representative example from the early 2000s: two major technology companies — fierce competitors, deeply distrustful of each other — negotiated a bilateral agreement to collaborate on a series of XML-based interoperability specifications. Each side was concerned that the other would use the collaboration to gain an unfair advantage.

The result was an agreement that ran over thirty pages and took more than a year to negotiate. It covered governance, IP, confidentiality, decision-making, dispute resolution, and every conceivable contingency — including elaborate mechanisms for what happened if the parties couldn't agree, involving written notices, escalation timelines, and fallback procedures so convoluted that the drafters themselves had difficulty explaining them. The confidentiality terms were unusually restrictive. The dispute resolution provisions reflected legal teams that had been in litigation with each other and expected to be again.

This was the state of the art when you wanted to do a standards collaboration. You pulled out this kind of agreement, and you spent six to nine months negotiating every detail with counterparts who assumed the worst about your intentions. It worked — in the sense that it produced specs — but the overhead was enormous, and the agreements were so complex that they became barriers to participation. Nobody new wanted to sign something they couldn't understand, and nobody who was already in wanted to try to amend it.

The agreements were also shaped by the relationships — and the distrust — between the parties. Language that seems bizarre in isolation often made sense as a response to a specific fear or a specific negotiation dynamic. But that context was lost once the agreement left the room. Future participants would read these clauses and have no idea why they existed. The language would migrate from one agreement to the next — copied, not understood — and accumulate like geological layers of prior disputes.

The next evolution was the form consortium agreement — a template that reduced the negotiation from a blank page to a starting point. A well-designed form might be 13 pages rather than 33, with modular IP provisions that could be swapped depending on the engagement. This cut setup time significantly, but still required negotiation and execution by every participant.

The current generation — represented by frameworks like JDF and the Community Specification License — eliminates the negotiation entirely. The governance terms are predefined and non-negotiable. The IP modes are selected from a menu. A new project can be chartered and operational in days. The tradeoff is flexibility — you can't customize the framework — but for most engagements, the predefined options fall within the range of reasonable, and the speed advantage is decisive.

Chartering Working Groups

When a new working group is formed, the charter is the foundational document. It defines:

Scope — what the group can and cannot work on. As discussed in Chapters 9 and 10, getting scope right is one of the hardest and most consequential decisions. Too narrow and the group can't adapt. Too broad and participants face open-ended patent exposure.

IPR mode — which patent policy applies. In organizations that offer multiple modes, this is selected at the working group level, not the organization level. Different working groups within the same organization can operate under different IP terms.

Deliverables — what the group is expected to produce. A specification? A reference implementation? A test suite? A profile? The deliverables determine the scope of the patent commitment and the criteria for declaring the work "done."

Decision-making rules — how the group reaches consensus, whether it uses voting, and what the thresholds are. These are covered in depth in Chapter 12.

The charter is often the last thing people want to spend time on — everyone is eager to start the technical work. But time spent on the charter is time saved later, because charter ambiguity becomes governance disputes once the work is underway.

The Pre-Draft Advantage

One governance dynamic that's rarely written down but profoundly affects outcomes: whoever writes the first draft has an outsized influence on the final result.

In most working groups, the initial draft specification — sometimes called the "seed" or "input document" — comes from one or two companies that have already built something and want to standardize it. The working group then iterates on that draft. In theory, the process is open and every participant can reshape the document. In practice, the initial draft sets the architecture, the terminology, the approach. Subsequent contributions tend to modify and extend rather than replace.

This matters for two reasons. First, the company whose technology starts as the initial draft has a significant backwards-compatibility advantage. Their existing products are likely to be aligned with the standard as it evolves. Competitors may need to adapt their implementations.

Second, the initial draft sets the frame for the patent commitment. The scope of the specification — and therefore the scope of the necessary claims — is shaped by the initial architecture. If your technology is the starting point, your patents are more likely to read on the final spec, but you're also more likely to have anticipated and managed that exposure.

Getting a seat at the table early — and ideally being the one who writes the first draft — is one of the most effective strategic moves in standards engagement.

Feedback Mechanisms and External Participation

Most standards organizations have mechanisms for people outside the membership to provide feedback — feedback agreements, public comment periods, liaison relationships with other organizations.

The lesson for practitioners: if your client is contributing through a feedback mechanism rather than as a full member, make sure they understand what rights they're granting. The feedback agreement may be more expansive than the membership agreement in some respects — particularly around patent commitments to contributed material.

11.4 The Linux Foundation Model

The Linux Foundation deserves specific discussion because of its scale and influence in the current landscape.

Operational Backbone

The Linux Foundation provides operational infrastructure for hundreds of projects — accounting, tax filing, bank accounts, event management, legal support, and program management. For projects hosted under the Linux Foundation, this eliminates the need to build operational capability from scratch. You don't need to file for nonprofit status, set up a bank account, or hire an accountant. The infrastructure exists.

This is a genuine advantage, particularly for projects that need to move quickly or don't have the resources to stand up an independent organization. The Linux Foundation's scale means it has established processes for everything from membership invoicing to conference logistics.

Community Governance vs. Corporate Governance

A dynamic that surfaces in many hosted projects — at the Linux Foundation and elsewhere — is the relationship between community governance and corporate governance.

The project's technical community — the engineers writing the spec or the code — operates by consensus, with influence earned through contribution and engagement. This is community governance. It's meritocratic in principle, even if it's imperfect in practice.

The project's organizational governance — the board, the membership tiers, the budget — operates by corporate governance rules. Board seats go to companies that pay the highest dues. Votes are cast by corporate representatives, not individual contributors.

These two governance models coexist within the same project, and they don't always agree. The technical community may reach consensus on a direction that the board doesn't support. The board may set priorities that the technical community doesn't share. Managing this tension is one of the most important — and most underappreciated — aspects of standards governance.

The best organizations acknowledge both models explicitly. They define which decisions belong to the technical community (spec content, architecture, feature prioritization) and which belong to the corporate governance (budget, membership, organizational strategy). When the boundary is unclear — as it often is — they have a process for resolving the conflict that doesn't require one side to simply override the other.

11.5 Competitive Dynamics in Governance

Standards governance exists in a competitive environment. The participants are often competitors, and the governance structure is a playing field on which competitive dynamics play out — sometimes constructively, sometimes not.

Constructive Competition

In a well-functioning standards body, competition drives better outcomes. Companies advocate for their technology, argue for approaches that align with their products, and push for specs that serve their customers. This is the intended dynamic. The process aggregates these competing interests and produces a specification that reflects some balance among them.

Destructive Competition

Sometimes participants engage for the purpose of slowing things down — not to improve the spec, but to prevent it from succeeding. A dominant incumbent may join a competing standards effort specifically to delay it, protecting their existing product from disruption. They'll propose endless amendments, argue every procedural point, and object to consensus without offering alternatives.

This is harder to detect and harder to counter than it sounds. In open source, you can overpower an obstructionist by committing more code. In standards, due process gives every participant a voice, and someone who contributes nothing technically can still exercise procedural rights. A small company that knows how to play the game can be remarkably effective at blocking progress — for good reasons or bad ones.

The flip side is also true. Standards give smaller companies power they wouldn't have in a purely market-driven competition. One vote per company, regardless of size, means that a startup with a good technical argument can prevail against much larger competitors. The process rewards preparation, persuasion, and persistence rather than raw market power.

The Substance-and-Process Dynamic

A pattern you'll see repeatedly: if you can't win on substance, you win on process. If you can't win on process, you win on substance. The most effective participants understand both and move fluidly between them. They know the technical arguments, but they also know the rules — voting thresholds, procedural requirements, charter limitations — and they use them strategically.

This is not cynicism. It's the reality of multi-party governance. The rules exist to create fair process. Understanding them well enough to use them effectively is part of being a competent participant.

This puts a premium on actually knowing the rules, especially in organizations that run on *Robert's Rules of Order* — and a surprising number do, including many ANSI-accredited bodies and most national-body delegations. *Robert's Rules* is several hundred pages of procedural detail, and most participants in the room have, at best, a vague memory of it from a high school student council. The participant who actually knows the rules — when a motion is debatable, when a point of order can be raised, what counts as a quorum, which supermajority threshold applies to which kind of motion, how to table versus postpone, how to call the question — has a structural advantage that has nothing to do with the merits of their position. He who knows the rules best wins. This is not an argument for studying *Robert's Rules of Order Newly Revised* cover to cover. It is an argument for knowing the procedural environment of every organization you operate in well enough that you don't get out-maneuvered by someone who does. At minimum, identify the one or two participants in each working group who actually know the rules. Either work with them or be ready for them.

11.6 Institutional Inertia and Organizational Evolution

Standards organizations, like all institutions, develop inertia. The rules that made sense when the organization was founded may not make sense ten years later. The membership that was appropriate for the initial scope may not be appropriate as the work evolves. The governance structure that worked for five companies may not work for fifty.

The challenge is that changing governance in a standards organization is hard. In a corporate entity, you amend the bylaws — a process that's well understood and has established legal mechanics. In a contractual consortium, you need every party to re-execute the agreement, which effectively gives everyone a veto. Even in organizations with formal amendment procedures, the political cost of proposing changes can deter anyone from trying.

The result is that many organizations operate under governance structures that no longer fit their reality. Working groups that have outgrown their charters. Membership tiers designed for a different era. Decision-making rules that were calibrated for a smaller group. The governance debt accumulates silently until a crisis forces the conversation.

One principle that helps: design for evolution from the beginning. Build amendment procedures that don't require unanimity. Use frameworks like JDF that allow working groups to be chartered and rechartered without renegotiating the organizational agreement. Separate the organizational governance (which should be stable) from the working group governance (which should be adaptable). The organizations that handle growth well are the ones that anticipated it.

And if the organization can't evolve, be prepared for the work to move elsewhere. The best way to prevent a fork is to make one unnecessary — by keeping the governance flexible enough to accommodate the community's changing needs. Or, as a practical matter: the best way to avoid forking is to allow forking. The threat of exit is what keeps institutions responsive.

The Longevity Problem

A related challenge is that standards organizations rarely die cleanly. Contractual consortia lose momentum as participants move on — but the agreement just sits there, and the spec remains in the field. Nobody formally winds down the organization. Nobody maintains the website. Ten years later, someone wants to implement the spec and there's no authoritative source, no point of contact, and no clear picture of what patent commitments are still in effect.

Incorporated entities have slightly better mechanics for winding down, but even there, dissolution is expensive and time-consuming. The more common outcome is a slow fade — meetings get less frequent, participation drops, the staff (if there was any) moves on, and the organization exists in name only.

For practitioners, this means thinking about longevity at the design stage. Where will the canonical spec live if the organization ceases to function? Are patent commitments durable beyond the organization's active life? Is there a successor organization or a repository that will maintain the materials? These questions feel premature when an organization is just getting started, but they're much harder to answer after the fact.

Modern approaches — hosting specs on GitHub, using durable frameworks like JDF that persist under the Linux Foundation's infrastructure — reduce this risk. But they don't eliminate it entirely. The governance structure should contemplate what happens when the work is done and the participants lose interest. Because they will.

Practice Tips

1. Determine whether patent commitments apply at the organization or working-group level.
2. When evaluating membership tiers, ask: "What are the patent obligations for this tier, and what rights do I receive?"
3. Whoever writes the initial draft of a spec has outsized influence on the final architecture.
4. Design for evolution: build amendment procedures that don't require unanimity.
5. Think about longevity: where will the canonical spec live if the organization fades?

Chapter 12 — Decision-Making and Voting

How a standards body reaches decisions is as important as what it decides. The decision-making structure determines who has power, how disagreements get resolved, and whether the resulting specification will actually be implemented. A standard that passes a vote but doesn't reflect genuine consensus is a standard that sits on a shelf.

One thing to understand from the outset: every organization has its own rules, processes, and culture around decision-making. There is no universal standard for standards governance. What constitutes a quorum at W3C is different from what constitutes a quorum at OASIS. How IETF gauges rough consensus is nothing like how ISO/IEC conducts a national body ballot. The voting thresholds, the role of the chair, the treatment of abstentions, and even the definition of "consensus" vary from organization to organization. When you engage with a new standards body, read the governance rules before you attend your first meeting. The assumptions you bring from your last engagement may not apply.

This chapter covers the mindset required for multi-party decision-making, the mechanics of consensus and voting, the group dynamics that shape outcomes, and what happens when consensus fails. The governance rules discussed here exist within an antitrust framework — standards bodies are fundamentally a creature of antitrust law, and the openness, balance, and due process requirements that shape governance design are what keep these competitor collaborations legal. That framework is covered in depth in Chapter 4.

12.1 The Mindset Shift

Multi-party negotiations require a fundamentally different approach than bilateral ones, and standards are multi-party negotiations at their most complex. The adversarial instincts, the advocacy skills, the habit of framing everything as win-or-lose — none of it works when the "other side" is fifteen companies that all need to agree and then voluntarily build to the result.

The core problem is structural. In litigation, a judge decides. In bilateral negotiation, you and a counterpart hammer it out, and if you can't agree, you walk away. In both contexts, there's a resolution mechanism that doesn't require the other side's voluntary cooperation after the fact.

Standards have no such mechanism. There's no judge. There's no binding resolution. And the output is voluntary — nobody is required to implement. You can win every vote and still lose if the companies that matter choose not to adopt the result. This changes the entire calculation. The goal isn't to prevail over the other side. It's to produce an outcome that enough participants will voluntarily implement to make the standard successful.

This requires a fundamentally different toolkit. Instead of advocacy, you need coalition-building. Instead of clever arguments, you need an understanding of what every participant needs and what they can live with. Instead of going for the win, you need to find the compromise that keeps enough parties at the table. The participants you're working with today are the same ones you'll encounter next quarter in a different working group and next year in a different organization. The relationships compound — for good or for bad.

In some international settings, the dynamic is literally diplomatic. UN-chartered organizations like ITU-T begin meetings with formal acknowledgments of national delegations. The pace is deliberate. The language is careful. The back-channel conversations are where the real work happens. Even in less formal settings — a consortium working group, a JDF project — the underlying dynamic is the same. You're building coalitions, managing relationships, and trading concessions across multiple dimensions simultaneously.

12.2 Consensus-Based Decision Making

What Consensus Actually Means

Consensus doesn't mean unanimity. This is one of the most common misunderstandings in standards governance.

Unanimity means everyone agrees. Consensus means the group has reached a position that everyone can live with — even if not everyone loves it. There may be participants who would have preferred a different outcome but aren't willing to block progress over it. Consensus is often best understood as the **absence of sustained and vigorous objection**, not the presence of universal enthusiasm.

This distinction is more subtle than it appears. Consider a room of ten people. If six support a proposal and four are ambivalent — they don't love it but they're not going to fight it — you likely have consensus. But if eight people enthusiastically support the proposal and two are vigorously and persistently opposed, you may not have consensus, even though you have a clear numerical majority. The intensity of the objection matters as much as the count.

In practice, the chair of a working group typically gauges consensus by asking whether anyone objects to a proposed direction. If no one objects, consensus is declared. If someone does object, the chair evaluates whether the objection is substantive, whether it represents a broader concern or an isolated position, and whether the objector is willing to stand aside or intends to sustain the objection. A single objection doesn't necessarily break consensus — but a sustained, vigorous one from even a small minority can.

The Difference Between Unanimity, Consensus, and Majority Rule

These three decision-making models exist on a spectrum, and different organizations use different models for different types of decisions.

Unanimity is rare in standards because it gives every participant a veto. It doesn't scale — in a large group, a single holdout can paralyze the process. But for small groups of five companies or fewer, you should generally strive for unanimous agreement even if the governance rules provide otherwise. In a group that small, every participant matters for adoption. Forcing a decision over the objection of one out of five participants is a hollow victory — because the "losing" party doesn't have to implement, and without their implementation, the standard may not achieve the critical mass it needs.

Consensus is used by many standards bodies, particularly in the technology space. It allows progress without requiring every participant to affirmatively agree, while still ensuring that serious objections are heard and addressed. The challenge is that consensus is inherently subjective — reasonable people can disagree about whether consensus has been reached, and the chair's judgment call is sometimes contested.

Majority voting is the default decision-making mechanism in many standards organizations, particularly incorporated entities with formal bylaws. It's clean and unambiguous — you count the votes, and the result is the result. The risk is that a majority can push through a decision over the objection of a significant minority, which may undermine the legitimacy of the standard and discourage implementation by the dissenting parties.

Some governance structures combine both approaches. The JDF model, for instance, uses consensus as the primary mechanism with a supermajority vote as a backstop, as discussed below. Other organizations use majority voting for most decisions but require supermajority or consensus for specific categories like IP policy changes.

The Consensus-Then-Supermajority Model

The JDF model illustrates a particularly effective approach: decisions are made by consensus, but if consensus can't be reached, the matter escalates to a supermajority vote.

The design is intentional. The supermajority threshold is high enough that neither side is likely to achieve it unless they have overwhelming support. This means that in most cases, the threat of a vote that nobody can win drives the parties back to the consensus process. The supermajority vote isn't really designed to produce a winner — it's designed to make both sides realize that they need to compromise, because neither can force the outcome through voting alone.

This is a governance structure that uses its own impracticality as a feature. The harder it is to win a vote, the stronger the incentive to find consensus. The supermajority backstop exists not to be used, but to make the consensus process work better.

Rough Consensus and Running Code

IETF's formulation — "rough consensus and running code" — deserves specific mention because it captures a philosophy that extends well beyond IETF.

Rough consensus acknowledges that perfect agreement is neither achievable nor necessary. What matters is that the group has broadly converged, that dissenting views have been heard, and that the remaining disagreements aren't fundamental. "Running code" means that the consensus is validated by implementation — a standard that works in practice is more valuable than one that's theoretically perfect but unimplemented.

This philosophy is specific to IETF's standards development approach and reflects its engineering culture. It is not how open source development works, and it is not the decision-making model used by most other standards bodies. But the underlying principle — that practical convergence matters more than theoretical perfection — is applicable broadly.

12.3 Voting Structures

When formal voting is required, the design of the voting structure shapes the outcome as much as the substance of the proposal.

One Vote Per Company

The most common structure in industry consortia is one vote per company, regardless of size. This prevents dominance by large companies — a startup's vote counts the same as the largest technology company in the world. It rewards persuasion and coalition-building over market power.

This principle typically extends to corporate families: a parent company and its affiliates collectively get one vote, not one vote per subsidiary. The rationale is straightforward — without this rule, a company could multiply its voting power by participating through multiple subsidiaries. The same logic that prevents one company from dominating through market size prevents it from dominating through corporate structure.

The limitation is that one-vote-per-company can also give disproportionate influence to small companies with narrow interests. A company with minimal implementation stake can cast a vote with the same weight as a company that will spend millions implementing the result.

Weighted Voting and Qualified Majority

Some organizations use weighted voting, where votes are allocated based on membership tier, contribution level, or other criteria. International organizations sometimes weight by national delegation. ETSI historically gave votes by national delegation, which — as discussed in Chapter 1 — led to companies joining multiple national delegations to multiply their voting power.

Qualified majority thresholds — requiring more than a simple majority to pass — are used for significant decisions like approving final specifications or amending governance documents. Common thresholds include two-thirds and three-quarters, though the specific number varies by organization and by the type of decision.

The Details Matter: Supermajority Precision

Supermajority thresholds sound straightforward until you try to apply them. A "two-thirds majority" can mean different things depending on the denominator.

Is it two-thirds of all members? Two-thirds of those present? Two-thirds of those who actually cast a vote — excluding abstentions? Each definition produces a different number, and the difference can determine the outcome on a close vote.

Even the arithmetic can become a question. Is two-thirds 66.6% or 66.7%? If you have 10 members voting and need a two-thirds majority, is that 7 votes or 6.67 (rounded to 7)? What if the answer changes with one abstention? These seem like pedantic questions until they decide a contentious vote, at which point the party that lost will scrutinize every decimal.

Different organizations answer these questions differently, and some don't answer them at all — they have thresholds defined in their governance documents without specifying the calculation methodology. When you're drafting governance rules, be precise. Specify the denominator. Specify how abstentions are treated. Specify the rounding. If the rules are ambiguous, the ambiguity will surface at the moment of maximum contention, which is the worst time to resolve it.

Organizations also vary in what decisions require a supermajority. Some require a supermajority only for governance changes (amending bylaws, admitting new membership classes) while using simple majority for technical decisions. Others require a supermajority for approving final specifications, on the theory that a standard that passes over significant opposition is unlikely to achieve broad implementation. A few require consensus for everything, with voting only as a last resort. The design choices here reflect the organization's priorities — speed versus inclusiveness, efficiency versus legitimacy.

Provisions That Require Unanimity — Or Can't Change at All

Some matters are so fundamental that they require unanimous agreement to change — or are designed not to change at all.

Patent policies are the primary example. In most well-designed organizations, changes to the IPR policy require unanimous consent of the affected members, or are simply prohibited. The rationale is that participants joined the organization based on the IP terms in effect at the time. Changing those terms after the fact — even by a supermajority — could alter the patent commitments participants have already made or alter the rights implementers are relying on.

This is why getting the patent policy right from the beginning matters so much. You can usually amend bylaws, adjust membership criteria, restructure working groups, and change decision-making thresholds through the governance process. But the patent policy is, for practical purposes, permanent. Treat it accordingly.

Voting Order and Visibility

Whether votes are cast simultaneously or sequentially, and whether they're visible or secret, creates different strategic dynamics.

Simultaneous, secret voting minimizes strategic behavior. Everyone votes based on their own assessment without being influenced by others' positions. This is cleanest but can produce surprising results when the outcome doesn't reflect the back-channel conversations.

Sequential, visible voting — as seen in the DVD-CCA structure discussed in Chapter 11 — creates signaling dynamics. Early voters set the tone. Later voters can adjust their positions based on what they've seen. This can facilitate coalition-building, but it can also enable manipulation, strategic abstention, and games around who gets stuck being the "bad guy."

Recorded votes with public minutes fall somewhere in between. The vote happens simultaneously, but the individual positions are recorded and published — meaning participants know their vote will be visible after the fact, even if they don't know how others voted at the moment they cast their own.

The choice between these models depends on the trust level among participants, the contentiousness of the issues, and whether the goal is to reveal true preferences or to build consensus through visible alignment.

Abstention and Quorum

Two mechanical issues that trip up organizations more often than you'd expect: abstention rules and quorum requirements.

Can a participant abstain? Does an abstention count toward quorum? Does it count as a "no" vote for purposes of calculating thresholds? Does it count as a vote at all for purposes of the denominator? Organizations answer these questions differently, and the answers can determine outcomes. An abstention that counts toward quorum but not toward the vote threshold lowers the effective bar for passage. An abstention that doesn't count toward quorum can prevent the vote from happening at all.

Quorum requirements — the minimum number of participants who must be present for a vote to be valid — prevent a small subgroup from making decisions for the whole. But they can also be used tactically: if you want to block a vote, simply don't show up and hope enough others don't either. Some organizations address this by allowing proxy voting or mail ballots to make quorum easier to achieve.

12.4 Group Dynamics at Scale

Small Groups vs. Large Organizations

Decision-making dynamics change fundamentally as groups grow.

In a group of three to five companies, you can work things out conversationally. Everyone knows everyone. Trust is personal. Consensus is reached through discussion, and formal voting is rarely needed.

At ten to fifteen companies, you start needing structure — formal agendas, defined decision-making rules, roles like chairs and editors. The personal dynamics still matter, but they're supplemented by process.

At fifty companies or more, the dynamics shift again. You can no longer negotiate with everyone individually. Coalition management becomes essential. The formal rules matter more because not everyone has a personal relationship with the chair. And paradoxically, it can be easier to drive decisions through a large group than a small one — because in a large group, most participants are deferential to the process and will go along with a well-presented proposal, while in a small group, every participant feels entitled to litigate every point.

Who Shows Up Matters

Standards decisions are made by the people who show up. This sounds obvious, but its implications are significant.

Attendance at working group meetings fluctuates. Some companies send the same representative every time. Others rotate, losing institutional context with each change. Some stop showing up when the work gets boring or the timeline slips.

The result is that the composition of the room can differ significantly from meeting to meeting, and the direction of the work can shift based on who happens to be present when a key decision is made. Companies that show up consistently — even when the agenda seems routine — have disproportionate influence on the outcome.

Persistent Objectors

Every standards body eventually encounters a participant who objects to everything. Sometimes the objections are substantive — the participant has a genuinely different technical vision. Sometimes they're strategic — the participant is trying to slow progress. Sometimes they're just temperamental.

The governance structure needs a way to handle persistent objection without either capitulating to it or ignoring it. Most organizations allow the chair to override a persistent objection after documented attempts to address the concern, subject to an appeal process. The appeal provides a safety valve — the persistent objector isn't silenced, but they can't hold the entire group hostage indefinitely.

12.5 Organizational Culture and Who's in the Room

The Evolution from Engineers to Governance Professionals

When organizations are small — five or ten companies in a working group — the people writing the specs are the same people making governance decisions. The chair is usually the most technically engaged person. The culture is engineering-driven. Governance is informal because everyone knows the rules, or there aren't many rules to know.

As organizations grow, a division emerges. The technical work stays with the engineers who build the specs. But the governance — the bylaws, the membership rules, the voting procedures, the IP policy — increasingly attracts a different type of participant: the standards professional. These are people whose job is governance, process, and organizational management. They may have engineering backgrounds, but their day-to-day work is running the organization rather than writing specifications.

This creates a cultural tension. The engineers want to build things. The governance professionals want to refine the rules. And governance professionals, in particular, have a tendency to want to "fix" governance — to propose amendments, restructure committees, and optimize processes. This can be productive when the governance genuinely needs improvement. It can also become self-referential work that consumes time and attention without advancing the technical mission. Be aware of this dynamic. Not every governance improvement proposal deserves the working group's time, and sometimes the best governance is the governance that stays out of the way.

Which Department Shows Up

Beyond the individual, pay attention to which department within a company sends the representative. The department's mission shapes what that representative optimizes for, and the results can be dramatically different.

In one content protection standards group, the major studios initially sent their content protection teams — the people whose job was to prevent piracy. These teams designed a system that was technically robust and extraordinarily complex. It maximized security. It addressed every conceivable piracy vector. Support costs, consumer experience, implementation burden — none of these were their problem. Their metric was content protection, and by that metric, the system was a success.

Later, the studios replaced their content protection representatives with people from their home video divisions — the people whose job was to sell movies to consumers through online channels. The home video teams walked into the working group and couldn't understand why the system was so difficult to implement, why the consumer experience was so poor, and why support costs were so high. After all, when a customer called because the movie they bought wouldn't play, it was the home video division that fielded the call and absorbed the cost.

Same companies. Same standards body. Entirely different priorities, because the representatives came from different departments with different success metrics. The lesson for practitioners: when you're evaluating the dynamics of a working group, don't just look at which companies are at the table. Look at who within those companies is sitting in the chair. Their department's mission will shape the standard.

12.6 When Consensus Fails

Sometimes the group can't agree. The technical visions are incompatible. The commercial interests diverge. The trust breaks down. What happens next depends on the governance structure and the stakes.

Appeals and Escalation

Most formal organizations have an appeal process — a participant who believes the chair declared consensus incorrectly, or that the process was unfair, can escalate to a higher body. The appeal mechanism is important because it provides legitimacy. Even if the appeal is denied, the fact that it was heard and considered gives the dissenter a path that doesn't require blowing up the process.

Fork Threats

The ultimate escalation in standards — as in open source — is the fork. If a group of participants is sufficiently unhappy with the direction, they can take the work (subject to copyright and IP terms) and start a competing effort elsewhere.

Fork threats are sometimes explicit and sometimes implicit. The HTML5/WHATWG story discussed in Chapter 1 is the canonical example — participants who disagreed with W3C's direction formed a competing group and ultimately won. The threat of a fork is what keeps standards bodies responsive to their participants. If the governance is too rigid, the process too slow, or the outcome too dominated by a single interest, the dissatisfied parties will leave.

This is why, as a governance design principle, the best way to avoid a fork is to make one unnecessary. Flexible governance, fair process, and genuine responsiveness to minority concerns reduce the incentive to walk away. The organizations that thrive long-term are the ones that give participants enough voice that they'd rather fix the problem from inside than start over from outside.

Walking Away

Sometimes the right answer is to accept that consensus isn't possible and move on. Not every collaboration succeeds. Not every standard gets adopted. The governance structure should contemplate this possibility — with clear rules about what happens to the IP, the materials, and the commitments when a working group concludes without a final deliverable.

The worst outcome is a zombie working group that never formally ends but never produces a result — consuming time, creating ongoing obligations, and preventing the participants from moving the work elsewhere. Clean termination provisions, with clear IP consequences, are worth the effort of negotiating up front.

Practice Tips

1. Read the governance rules before attending your first meeting — every organization's rules are different.
2. When evaluating voting thresholds, be precise about the denominator and how abstentions are counted.
3. For small groups of five or fewer, strive for unanimity even if the rules don't require it.
4. Identify which decisions require supermajority, which require simple majority, and which require consensus.
5. Pay attention to which department sends the representative — their success metrics shape the standard.
6. Consistent attendance matters more than company size.

Chapter 13 — Specification Development and Finalization

This chapter covers the practical mechanics of getting from a blank page to a published specification — the lifecycle, the key distinctions in spec drafting, compliance programs, versioning, and the path to international recognition.

13.1 The Lifecycle of a Specification

Pre-Draft and the Power of the First Proposal

Every specification starts somewhere, and where it starts matters more than most people realize.

In some working groups, the initial draft emerges organically from group discussion — everyone sits down with a clean sheet of paper and drafts together. But more often, someone shows up with an existing proposal. A company that has technology in the market, a research team with a prototype, or an engineer with a detailed vision contributes an initial document that becomes the starting point for the group's work.

As discussed in Chapter 11, this **pre-draft advantage** is one of the most powerful dynamics in standards. The initial architecture, the terminology, the technical approach — these all carry forward into the final spec, even after extensive iteration. A company whose technology starts as the pre-draft has a significant backwards-compatibility advantage. Their existing products are more likely to align with the final standard.

In organizations like JDF, there's a formal concept of a "pre-draft" submission. Multiple companies may bring competing proposals, and the group has to decide which one to use as the starting point. This decision is often political as much as technical. If there are three competing proposals and the group takes a majority vote, the losers may simply leave — and their implementation support leaves with them. The better approach is usually to combine proposals where possible, even if the result is an awkward starting point, because the goal is to keep everyone at the table.

Working Drafts and Iteration

Once the starting point is established, the work proceeds through iterative drafts. Engineers propose changes, the group discusses them, and the chair incorporates what has consensus. This is where most of the technical work happens — draft by draft, issue by issue.

From an IP perspective, this phase is important but often overlooked. Patent commitments typically don't lock in until the spec is finalized. Contributions during the drafting phase are provisional. This means participants can evaluate the evolving spec and decide whether to stay or withdraw before the commitment crystallizes.

One practical note: keep track of what's changing between drafts. If the scope drifts — the microwave starts becoming a toaster — that's when you need to flag the issue, not after the final draft is published.

Public Review and Comment Resolution

Most standards processes include a public review period before finalization. The draft is published, and parties outside the working group — including potential implementers, competitors, and the broader community — have an opportunity to comment. ANSI-accredited organizations are required to provide a public review period as part of their due process obligations, and many non-accredited organizations follow similar practices because it strengthens the legitimacy of the standard.

Comment resolution is one of the most process-intensive parts of standards development. Every comment must be reviewed. Every objection must be addressed — not necessarily accepted, but responded to with documented rationale. Organizations that skip this step, or that handle it pro forma, create governance risk. As discussed in Chapter 4, the right to be heard is a core element of due process, and documented comment resolution is part of the antitrust defense.

Final Approval and Publication

The final step is formal approval by the working group (and sometimes by a higher governing body) followed by publication. This is when patent commitments lock in, exclusion windows close, and the spec becomes the canonical reference that implementers build to.

The publication format matters. Some organizations sell the spec — ISO/IEC standards, for instance, are available for purchase. Others publish freely. Some use Creative Commons licenses. Some put specs on wikis. As discussed in Chapter 3, the copyright and distribution terms should support the goals of the standard: broad access for implementation, protection against unauthorized forking, and a canonical version that implementers can rely on.

13.2 Normative vs. Informative Text

Every specification contains two types of content, and the distinction has direct consequences for patent commitments.

Normative text is the part of the spec that implementations must follow. It describes requirements, interfaces, and behaviors that are necessary for compliance. Normative text is what patent commitments cover — if a claim reads on normative text, it may be a necessary claim.

Informative text — examples, explanations, appendices marked as non-normative, implementation guidance — exists to help readers understand the spec but doesn't create compliance requirements. Patent commitments generally don't extend to informative text.

This distinction matters because spec authors sometimes put substantive content in informative sections, either inadvertently or to reduce the patent commitment surface area. If a critical feature is described only in an informative appendix, implementers who rely on it may not have patent protection. When reviewing a spec, check where the technically important content lives. If it's informative, ask whether it should be normative — and understand the IP consequences of the answer.

Reference Implementations and Test Suites

Some standards bodies produce reference implementations alongside their specs — working code that demonstrates the spec can be implemented and that serves as a testable baseline.

Traditionally, reference implementations were intended solely for verification: does the spec work? Can it be implemented? If the reference implementation doesn't work, is the problem in the code or in the spec? These implementations were not intended for production use and were typically unoptimized.

Increasingly, standards bodies are producing open source implementations intended for actual deployment. This creates a convergence with the open source world discussed throughout this book — and it creates governance questions. Is the code subject to the standards body's patent policy or to its open source license? What happens when the open source implementation diverges from the spec? Who is the authoritative source — the document or the code?

Test suites serve a related function. They provide automated verification that an implementation conforms to the spec. Test suites are particularly valuable for interoperability — two implementations that both pass the test suite should interoperate, at least for the features the suite covers.

When Code Is the Spec

In some domains — encryption algorithms, compression codecs, certain protocol behaviors — the specification cannot be fully expressed in human-readable text. The algorithm itself is the standard. In these cases, the code is not a reference implementation of the spec. It is the spec.

This has IP implications. If the code is the normative specification, then patents on the algorithm — not just on the interface — may be necessary claims. The usual distinction between interface patents (covered by the commitment) and implementation patents (not covered) breaks down when the implementation is the standard.

13.3 Compliance and Certification

Self-Certification vs. Third-Party Testing

As discussed in Chapter 3, conformance programs come in several forms. Self-certification is lightweight and common. Third-party testing is heavier but provides more assurance. Hybrid approaches — formal testing for the first product, self-certification for subsequent ones — balance rigor with efficiency.

The choice between these models depends on the stakes. For safety-critical standards, third-party testing is often required by regulation. For software interoperability standards, self-certification is usually sufficient.

Certification Marks and Branding

The certification mark — the logo that goes on the product — is often the most commercially significant piece of IP in the standards ecosystem. As discussed in Chapter 3, retailers may refuse to stock products without the logo, and customs officers may seize products that bear the logo without proper licensing.

For practitioners, the key question is what the certification program actually tests. If it tests only normative, mandatory features, it aligns with the spec. If it tests optional features — elevating them to de facto requirements through the certification process — there's a gap between what the spec says and what the market demands. That gap creates both compliance risk and patent risk.

Interoperability Testing Events

Plugfests — organized events where multiple companies bring their implementations and test interoperability in real time — are one of the most effective tools for ensuring a spec works in practice. They surface ambiguities in the spec, implementation bugs, and interoperability issues that no amount of document review can find.

From a legal perspective, plugfests require careful handling. Competitors are sharing technical details about their implementations in a common venue. Confidentiality rules need to be clear. Patent implications need to be understood. And the results — which products interoperate and which don't — can be commercially sensitive.

13.4 Versioning and Maintenance

Minor Revisions vs. Major Versions

Standards evolve. Bug fixes, clarifications, and minor enhancements produce point releases (1.0 → 1.1 → 1.2). Significant changes in scope, architecture, or functionality produce major versions (1.0 → 2.0).

The distinction matters for patent commitments. As discussed in Chapters 7 and 8, a commitment to version 1.0 may or may not carry forward to version 1.1 or 2.0, depending on the policy. Minor revisions that correct errors without changing the normative content are usually covered by the original commitment. Major versions that add new normative content may require new commitments — and trigger new exclusion windows.

Backward Compatibility

One of the most difficult design decisions in versioning is how to handle backward compatibility. A new version that breaks compatibility with the old version fragments the ecosystem. A new version that maintains perfect backward compatibility may be constrained in what it can improve.

The approach varies by domain. Web standards tend to prioritize backward compatibility heavily — the web can't break existing websites. Protocol standards for new industries may prioritize feature advancement and accept that older implementations will need updating.

Sunsetting and Deprecation

Standards rarely die formally. They fade — implementations persist in the field long after the working group has moved on, and the spec sits on a shelf (or a website) indefinitely. As discussed in Chapter 11, the longevity problem is one of the most under-appreciated challenges in standards governance.

Explicit deprecation policies — formal declarations that a version is no longer maintained, no longer recommended for new implementations, and will eventually lose support — help manage this lifecycle. But even deprecated standards can persist in the field for decades. The VHS tape is no longer manufactured, but VHS players still exist. Standards work the same way.

13.5 International Transposition

The PAS Process

As discussed in Chapter 2, the PAS (Publicly Available Specification) process allows a consortium-developed specification to be submitted to ISO/IEC JTC-1 for an international vote. If the national bodies approve, the consortium spec becomes a formal international standard — with all the treaty-level weight, government procurement advantages, and regulatory recognition that entails.

The PAS process is a bridge between the speed of consortium development and the formality of international recognition. You develop where you want, then layer the international status on afterward. Organizations like W3C and the Joint Development Foundation have PAS submitter status, which means their specs can go through this process.

There's also a strategic use of the PAS process that's worth naming because it drives a meaningful share of the submissions that actually go through. Industry sometimes pushes a specification up to ISO/IEC not because it needs the international imprimatur for its own sake, but because international standards transpose downward into national standards systems in jurisdictions where consortium specs do not. China is the canonical example. A consortium specification developed at, say, the USB Implementers Forum or the Trusted Computing Group has limited standing in the Chinese national standards system on its own. The same specification, once accepted as an ISO/IEC international standard, can be transposed to a GB national standard and recognized in Chinese procurement and regulation. USB, UPnP, and TPM specifications all followed roughly that path. When you see industry investing in a PAS submission for a spec that already has effectively universal market adoption, the answer to "why bother?" is often a downstream national adoption strategy in a market where the international route is the only practical one.

National Body Voting and Comment Resolution

The international ballot process is its own governance challenge. National bodies vote — not individual companies. Each country's position is determined by its national body through its own process. Companies lobby their national body for the vote they want, but ultimately the vote is the nation's, not the company's.

The comment resolution process at the international level is formal and documented. National bodies submit comments with their votes, and the submitting organization must address those comments before the spec can advance. This is where the diplomatic skills discussed in Chapter 14 become relevant — building coalitions across national bodies, understanding each country's concerns, and managing the process to achieve the required voting threshold.

Regional Regulatory Adoption

Beyond formal international standardization, standards are increasingly adopted by reference in regional regulations — the EU Cyber Resilience Act, AI regulations, data protection frameworks. When a regulation references a standard, compliance with the standard becomes a pathway to regulatory compliance. This gives the standard significant commercial value and creates incentives for participation in its development.

For practitioners, the implication is that standards work doesn't end at publication. Understanding the regulatory landscape — which standards are being referenced by which regulations in which jurisdictions — is increasingly part of the advisory role.

Practice Tips

1. Invest time in the charter — charter ambiguity becomes governance disputes once the work is underway.
2. Check where technically important content lives — if it's in informative sections, it may not be covered by patent commitments.
3. Understand whether the reference implementation is for verification only or intended for production deployment.
4. Monitor normative references — they import the referenced spec's IP terms into your specification.
5. Think about versioning and backward compatibility from the beginning, not after the first version ships.

Chapter 14 — The Art of Multi-Party Negotiation in Standards

Much of this book has focused on the legal frameworks — patent policies, governance structures, voting mechanics. This chapter is about the human side: how to actually get things done in a room full of competitors, each with their own agenda, their own constraints, and their own definition of success.

Standards negotiation is a distinct skill. It's not litigation. It's not bilateral deal-making. It's not open source community management. It borrows from all three but operates by its own rules. The practitioners who are most effective in this space understand those rules intuitively. This chapter tries to make them explicit.

14.1 The Fundamental Difference: It's Not Bilateral

Chapter 12 discussed the mindset shift required for standards work — the move from litigation thinking and bilateral negotiation toward diplomacy, empathy, and coalition management. This chapter builds on that foundation with the practical tactics of how to get things done in multi-party settings.

From Individuals to Coalitions

Standards decisions are made by coalitions, not by individuals. In most consortia, even the largest company in the room has one vote. To get your preferred outcome, you need others to support it — and ideally to advocate for it on your behalf so you don't have to.

The most effective approach is to do the groundwork before the meeting. Back-channel conversations. Informal alignment. Understanding where the votes are and where they aren't. If you've done this well, the formal meeting is a confirmation — the positions are already set, and the chair is declaring consensus that was built outside the room.

If you haven't done the groundwork, the meeting becomes unpredictable. Proposals that seemed reasonable get challenged. Votes that seemed assured are lost. Participants who seemed supportive abstain. The meeting is where the result is announced, not where it's created.

This is what makes standards more like diplomacy than law. In the UN, the important conversations happen in the hallways between sessions, not at the podium. Standards work the same way. The corridor conversations, the dinners, the informal calls between meetings — these are where alignment happens. The formal process ratifies what the informal process produced.

Identifying Trades

Once you understand what everyone needs (Chapter 12's empathy imperative), you can identify trades — places where you can support someone else's priority in exchange for support on yours.

The key is knowing what you truly care about and what you can concede. If you care deeply about the patent policy but are indifferent to the governance structure, you can trade your support on governance in exchange for someone else's support on IP. If another participant cares about certification criteria but not about scope, there's a natural trade.

These trades happen across issues within a single working group, but they also happen across working groups and even across organizations. A relationship built by supporting a company's position in one forum creates goodwill that carries into the next. The currency is cumulative.

Strategic Support and Reciprocity

Sometimes it makes sense to actively support another company on an issue they care deeply about — even if you don't care about it at all. Reciprocity is a deeply human instinct. If you throw your weight behind someone else's priority when it costs you nothing, they're far more likely to support you when you need it on something that does matter.

This is different from transactional deal-making, where you explicitly negotiate quid pro quo. In standards, the reciprocity is often informal and unspoken. You support a company's position in a governance discussion. Three months later, they support yours in a patent policy debate. Nobody said "I'll do this if you do that." It just works — because people remember who helped them and who didn't. You need friends before you need friends. Build the goodwill now, and it will be there when you need it.

This also means that repeat players have a built-in advantage. If you've been working in a community for years, you have relationships, credibility, and a track record of reciprocity. If you're new to the group, resist the urge to push hard on your first day. Read the room. Understand the existing dynamics and alliances. Concentrate on building relationships before you try to spend them. The investment pays off quickly — but only if you make it before you need the return.

It's all about the chits — the accumulated goodwill and trust you've built over time. If you want something, you need chits to spend. And if you're new, you don't have any yet. The only way to get them is to earn them — by showing up, by helping others, by being fair, and by doing it before you need something in return.

The corollary is equally important: **pick your fights wisely.** You can't win on everything. If you fight every battle with equal intensity, you dilute your influence and exhaust the goodwill of potential allies. The most effective participants identify the two or three issues that truly matter to their organization and focus their energy there — even if it means conceding on other points. Losing gracefully on something you don't care about preserves your credibility and your relationships for the fights that count.

The Power of "No"

One of the most effective negotiating tools in standards is the ability to say no — not "no, because," which invites people to solve your stated problem and come back with a revised proposal, but simply no. Once you explain your reasons, the other side will work to address them and return with something you may find harder to reject. Sometimes the strongest position is the unexplained one.

This doesn't mean being obstructionist. It means being selective about when to engage on substance and when to hold a position without elaboration. It's a tool to use sparingly, but when deployed correctly, it signals that this is a genuine boundary, not a negotiating position waiting for a better offer.

14.2 The Players and Their Motivations

Understanding the motivations of the participants is the foundation of effective standards negotiation. Different types of organizations engage for different reasons, operate under different constraints, and respond to different incentives.

Big Tech Companies and Platform Strategies

Large technology companies participate in standards for strategic reasons. A standard that aligns with their platform creates an ecosystem around their technology. A standard that disrupts their platform threatens their business. Understanding which dynamic is at play — and for which participant — is essential for predicting behavior.

Large companies also tend to have internal complexity that affects their standards engagement. The engineer in the working group may have a different view than the patent team, which may have a different view than the business unit, which may have a different view than the legal team. The company's "position" may not be as unified as it appears. This creates both risk and opportunity — risk that a commitment made by the engineer gets overridden later, and opportunity to find alignment with one faction even if another is opposed.

Watch for the internal dynamics. A company that sends its engineers tends to produce technically sound positions but may not have cleared the IP implications. A company that sends its standards diplomats tends to produce politically savvy positions but may not understand the technical tradeoffs. The most effective companies send both — and they coordinate.

Startups and VC-Funded Companies

Startups engage in standards for different reasons than large companies. They typically want to standardize their technology to create a larger market, to gain credibility with enterprise customers, or to prevent a larger competitor from controlling the interface.

Startups bring energy and technical depth. They also bring constraints: limited legal resources, pressure to move fast, and corporate structures that can create patent commitment complications. As discussed in Chapter 6, a VC-backed company controlled by a fund that also controls dozens of unrelated companies may have genuine difficulty making the affiliate commitments that patent policies require.

Don't underestimate startups in the standards process. A small company that knows how to play the game can punch well above its weight. Remember: one company, one vote, regardless of size. A startup with strong technical arguments, good relationships, and consistent attendance can be more effective than a large company that sends different people every time and can't keep its internal positions aligned. Some of the most effective standards participants I've encountered have been small companies with deep domain expertise and the discipline to show up every time.

Individual Contributors and Academics

Not everyone at the table represents a company. Some participants are independent consultants, academic researchers, or individual contributors who participate on their own behalf. These participants can be highly effective — they often bring deep technical knowledge, long institutional memory, and credibility that isn't tied to any company's commercial agenda.

They can also be unpredictable. Without a corporate principal to answer to, their positions may be driven by personal philosophy, academic interest, or long-standing views about how technology should work. Understanding their motivations requires more nuance than understanding corporate motivations, where the incentives are usually commercial and therefore predictable.

Complex Corporate Structures

Some participants come from organizations where the corporate structure itself creates complications. Conglomerates with multiple business units, each with their own P&L and their own patent portfolio. Subsidiaries that have no authority to bind the parent's patents. VC portfolio companies where the fund's control creates affiliate relationships with dozens of unrelated companies.

These structural issues affect negotiation dynamics beyond just patent commitments. A participant who can't commit their affiliate's patents has less to offer and less to trade. A participant whose internal approval process takes six months can't respond to proposals on the timeline the working group wants. Understanding these constraints — and accommodating them where possible — is part of effective multi-party management.

14.3 Influence Without Authority

Chair and Editor Roles

Beyond the governance functions discussed in Chapter 11, the chair and editor of a working group are **negotiation actors** with disproportionate influence on outcomes.

The chair controls the agenda — which issues get discussed, in what order, and for how long. They determine when consensus has been reached — a judgment call that reasonable people can disagree about. They manage the meeting dynamics — calling on speakers, cutting off unproductive debates, and steering the conversation toward resolution. A sympathetic chair who understands your priorities can create opportunities. An adversarial chair can systematically disadvantage your positions through agenda management and consensus calls.

The editor decides how consensus positions get translated into words. The choice of language, the structure of the document, the way ambiguities are resolved in drafting — these are all editorial decisions that shape the final spec in ways that may not be visible in the working group discussions. An editor who subtly favors one technical approach over another in how they draft normative text can steer the spec without the working group explicitly deciding.

For practitioners advising on a new engagement, the identity of the chair and editor is as important as the governance rules. If possible, have your people in those roles. If not, at least understand who holds them and what their tendencies are.

Technical Credibility as Currency

In standards, the most effective advocates are the ones with the deepest technical knowledge. A participant who can explain why a particular approach is better — with specifics, with data, with implementation experience — carries more weight than a participant who argues from commercial preference.

This is one of the ways standards differ from open source. In open source, you earn influence by writing code. In standards, you earn influence by writing specs, proposing solutions, and demonstrating that your technical judgment is sound. The currency is credibility — and once you have it, your positions carry weight even when the argument isn't airtight.

Credibility also means admitting when you're wrong. A participant who acknowledges a flaw in their proposal and offers to fix it earns more respect than one who digs in. Standards professionals have long memories, and a reputation for intellectual honesty is worth more than winning any single argument.

Securing the Pre-Draft

Building on Chapter 11's discussion of governance dynamics, the pre-draft advantage is one of the highest-impact negotiation strategies available. Whoever writes the first draft shapes the architecture, the terminology, and the technical direction of the final spec.

Securing the pre-draft position requires three things: **technical readiness** (having something concrete to propose), **political groundwork** (having enough support among other participants that your proposal gets accepted as the starting point), and **timing** (being ready before anyone else is). Companies that invest in developing a proposal before the working group even forms — and that line up support from key participants in advance — have a structural advantage that persists throughout the life of the project.

When multiple pre-drafts compete, the resolution is often a combined document — an awkward Frankenstein starting point that incorporates elements from each proposal. This is messy but functional. The alternative — a vote where the losers walk away — is worse. Getting an imperfect proposal accepted that keeps everyone at the table is better than winning a clean vote that sends half the participants home.

14.4 Tactical Considerations

Managing the Pace

One of the most underappreciated tactical levers in standards is pace. If you want a standard to succeed, push for faster timelines — shorter review periods, more frequent meetings, aggressive milestones. If you want to slow a competing standard or buy time for your own technology, advocate for longer review periods, more thorough due diligence, and deliberate process.

Both approaches are legitimate. Due process requires adequate time for review and comment, and there's always a genuine argument that the group should be more thorough. The skill is in knowing when thoroughness serves the standard and when it serves a competitive agenda — and being honest with yourself about which one you're pursuing.

When to Engage and When to Observe

Not every working group meeting requires active participation. Sometimes the most valuable thing you can do is observe — understand the dynamics, identify the key players, map the positions, and wait for the right moment to engage. Early engagement commits you to positions that may need to change as the work evolves. Late engagement lets you see where the spec is going before you invest.

The risk of waiting too long is that the architecture gets set without your input, and changing direction becomes increasingly expensive. The risk of engaging too early is that you burn political capital on positions that become irrelevant. The judgment call is context-dependent, but as a general principle: engage early on scope and architecture (which are hard to change later) and later on details (which are easier to adjust).

Bikeshedding — The Trap of Trivial Debates

In 1957, C. Northcote Parkinson — a British naval historian with a gift for organizational satire — published a book called *Parkinson's Law*. Most people know the title principle: work expands to fill the time available. But buried in the same book was a less famous observation that anyone who has sat through a standards meeting will recognize immediately.

Parkinson described a fictional finance committee reviewing three agenda items. The first was a proposal to build a £10 million nuclear reactor. The committee approved it in two and a half minutes. The technology was so complex that only two members understood it, and neither wanted to admit uncertainty by asking questions. Everyone else deferred.

The second item was a proposal to build a £350 bicycle shed for the staff. This unleashed a forty-five-minute debate. Everyone understood bicycle sheds. Everyone had an opinion on the materials, the design, the placement, the roof. One member suggested aluminum. Another preferred asbestos. A third questioned whether a shed was needed at all. The discussion consumed more time than the nuclear reactor and the third agenda item combined.

The third item — £21 for refreshments at a committee meeting — generated an even longer debate, because the amounts involved were within every member's personal experience. People who couldn't visualize ten million pounds had very strong feelings about the price of coffee.

Parkinson's point was that the amount of time spent on an agenda item is inversely proportional to its importance. The more complex and consequential the decision, the fewer people feel qualified to engage — and the faster it passes. The more trivial and accessible the decision, the more opinions flood in — and the longer it takes.

He was writing about government committees in the 1950s. He could have been writing about any standards body meeting in the 2020s.

Standards work is extraordinarily susceptible to this. The hard architectural decisions — scope, patent policy, normative structure — often pass with minimal debate because only a handful of people in the room truly understand the implications. A scoping decision that determines the outer boundary of every participant's patent commitment may get ten minutes of discussion. A debate over whether the spec should use "must" or "shall" in a particular clause — or whether a diagram should use blue or green arrows — can consume the rest of the afternoon. Everyone can participate. Everyone has a view. And nobody wants to concede on something they understand, precisely because they understand it.

The danger isn't just wasted time. It's that the real decisions get made while the room is arguing about the bike shed. The participants who recognize this dynamic focus their energy on the architectural questions and let the trivial debates run their course. The ones who don't end up spending their political capital on issues that won't matter six months from now.

This isn't unique to standards. Lawyers do it to themselves in every negotiation. Two legal teams will spend three rounds of redlines on choice of law, warranty disclaimers, and indemnification caps — provisions that are well-understood, that every attorney feels comfortable marking up, and that will almost certainly never be litigated. Meanwhile, the provisions that actually determine the economics of the deal — the scope of the license, the definition of the deliverables, the termination triggers — get a fraction of the attention because they're harder, less familiar, and require the lawyers to understand the business context rather than just the legal boilerplate. If you catch yourself deep in a markup of provisions you've negotiated a hundred times before, ask whether the reactor just passed while you were painting the bike shed.

But bikeshedding is also a tool — and a diagnostic.

As a tool, it can work to your advantage. If you need time — to align your internal stakeholders, to wait for a patent review, to let a competing proposal lose momentum — a well-placed trivial debate can absorb a meeting without touching the substance you're not ready to address. You don't have to start the bikeshed debate. You just have to not stop it. Let the room spend forty-five minutes on terminology while the architectural question you're not prepared to fight over yet sits safely on next month's agenda.

As a diagnostic, pay attention to who is bikeshedding and why. When a participant who normally engages on substance suddenly becomes passionate about a trivial issue, ask yourself what they're avoiding. Are they stalling? Are they uncomfortable with the direction of the real work and looking for a way to slow it down without objecting on substance? Or are they simply out of their depth on the technical questions and gravitating toward what they can contribute to?

Understanding the difference tells you a lot about what's actually happening in the room. A participant who bikesheds because they're stalling is playing a strategic game — and you should be asking what they're buying time for. A participant who bikesheds because the substance is over their head is telling you who their organization actually sent to the meeting, which may not be who you expected.

If you find yourself in a meeting where the group has spent thirty minutes debating something that has no IP, governance, or technical consequence — that's the bike shed. Decide whether to let it burn, use it, or shut it down. All three are valid choices depending on what you need.

Process as Toolkit: The Sabotage Manual Sitting in Plain Sight

In 1944, the Office of Strategic Services — the wartime predecessor to the CIA — distributed a small booklet called the *Simple Sabotage Field Manual*. It was meant for ordinary citizens in occupied territory: assembly-line workers, clerks, anyone embedded inside an enemy organization who couldn't risk armed resistance but could quietly degrade output. The most famous section is titled "General Interference with Organizations and Production." It reads like a transcript of a difficult standards meeting, and it was written by people who had never attended one.

A partial list of the recommended tactics:

- "Insist on doing everything through 'channels.' Never permit short-cuts to be taken in order to expedite decisions."
- "Make 'speeches.' Talk as frequently as possible and at great length."
- "Refer all matters to committees, for 'further study and consideration.'"
- "Bring up irrelevant issues as frequently as possible."
- "Haggle over precise wordings of communications, minutes, resolutions."
- "Refer back to matters decided upon at the last meeting and attempt to re-open the question of the advisability of that decision."

- "Advocate 'caution.' Be 'reasonable' and urge your fellow-conferees to be 'reasonable' and avoid haste which might result in embarrassments or difficulties later on."
- "Be worried about the propriety of any decision — raise the question of whether such action as is contemplated lies within the jurisdiction of the group."

Anyone who has spent a year inside an active working group has watched every one of those, often in the same meeting. The interesting thing isn't that the tactics exist — it's that they're indistinguishable from procedural diligence. Insisting on channels is what governance documents tell you to do. Re-opening last meeting's decision can be a perfectly legitimate request to revisit something that's no longer fit for purpose. Worrying about jurisdiction is exactly what working-group charters are supposed to discipline. The behaviors are the same whether they come from a sophisticated participant deliberately slowing the work, a junior delegate who has been told to be cautious and is doing their job, or a chair who genuinely thinks the group needs more time.

That's why the manual is useful. Three things to do with it.

First, distinguish deliberate use from well-meaning diligence. Both look identical in the moment, but they call for very different responses. A pattern of objections from a sophisticated repeat player who is also the one whose product would lose if the spec ships on time is probably strategic. The same objections from a newly-assigned attorney who has just read the bylaws for the first time and is genuinely worried about jurisdiction are not. Strategic process moves are best handled with counter-process: get the objection on the record, get an answer to it, get a decision, move on. Earnest process anxiety is best handled with patience and education — answer the actual question, walk through the actual rule, and the participant will usually relax once they see the work has the cover they were looking for.

Second, recognize the pattern even when no individual move looks suspicious. Each tactic in isolation reads as good citizenship. The signal is in the aggregation: the same participant raising fresh procedural concerns at every meeting, on different topics, with no underlying technical position. That's the diagnostic. You don't have to confront it — naming it for yourself or to the chair changes how you allocate your attention and your time.

Third, know when to use the toolkit yourself. There are moments when you need time. Your internal stakeholders haven't aligned. Your patent review hasn't come back. A competing proposal is losing momentum and another two weeks would let it die on its own. The same tools that look like sabotage in someone else's hands look like prudence in yours. A request for further study, a careful re-examination of jurisdiction, a thorough review of the minutes from the last meeting — all of these are legitimate, and all of them buy you the time the situation actually requires. The question to ask before reaching for one of them is whether what you're protecting is worth the small amount of institutional friction you'll create. Most of the time, used sparingly, it is. Used as a default, it isn't, because the other repeat players keep score.

The manual isn't a list of villainous practices. It's a list of moves the institution makes available to anyone who knows where to look. Knowing they're there — and recognizing them when others use them — is part of the practitioner's basic literacy.

Managing Internal Stakeholders

For attorneys advising on standards engagements, one of the most challenging aspects is managing the internal client. The engineer in the working group has their own view of the right technical direction. The patent team has concerns about exposure. The business unit wants the standard to align with their product roadmap. The legal team wants manageable risk.

These internal stakeholders often have conflicting priorities, and the attorney is sometimes the only person who sees all the dimensions. Managing these conflicts — ensuring the engineer doesn't make IP commitments the patent team hasn't approved, ensuring the patent team doesn't block engagement the business unit needs, ensuring the business unit understands the constraints of the standards process — is as much a part of the job as the external negotiation.

14.5 Playing the Long Game

Standards as Multi-Year Commitments

Most standards engagements run for years. The initial formation takes months. The technical work takes a year or more. The review, finalization, and adoption cycle adds another year. And then maintenance, versioning, and evolution continue indefinitely.

This timeline means that the participants you work with in year one are the same participants you'll work with in year three. The favors you do — or the enemies you make — compound over time. A reputation for fairness, competence, and good faith is built slowly and carries enormous value. A reputation for gamesmanship or bad faith is equally durable.

Reputation and Trust Across Organizations

Standards professionals tend to work across multiple organizations. The person who represents a company at W3C may also represent them at OASIS, at JDF, and in bilateral negotiations. The relationships you build in one venue carry into others.

This cross-pollination is one reason why trust is the most important currency in standards work. A promise kept in one organization builds credibility in the next. A commitment broken in one venue follows you everywhere. There's plenty of credit to go around in standards work — share it generously, and it comes back.

The Professional Participant

One dynamic worth noting: some participants are professionals whose primary job is standards engagement. They attend meetings across multiple organizations, build deep expertise in governance and process, and develop relationships that span decades.

These professional participants can be extraordinarily effective allies — they know the rules, they know the people, and they know how to navigate the process. They can also become institutional unto themselves, pursuing agendas that serve their personal standing in the standards community as much as their employer's commercial interests. As discussed in Chapter 2, people develop strong affinities for the organizations they know, and sometimes those affinities shape their recommendations.

When you encounter a professional participant, understand their motivations. They may be your most valuable resource — or the most sophisticated challenge you'll face. Either way, they're not going away.

Shape the Ecosystem Before You Need It

The most strategic thing you can do in standards is invest in relationships and governance structures before you have a specific need. Join organizations where your technology might eventually need a standard. Build relationships with the people who will be in the room when the time comes. Contribute to governance improvements that create a favorable environment for your future work.

If you wait until you need a standard to start engaging, you're already behind. The companies that are most effective in standards are the ones that shaped the ecosystem years before they needed it. They wrote the governance rules. They built the relationships. They established the trust. And when the time came to standardize something that mattered to them, they were playing on their own gameboard.

Practice Tips

1. Do the groundwork before the meeting — the meeting is where the result is announced, not where it's created.
2. Identify what you truly care about and what you can concede — then trade strategically.
3. Support others on issues they care about when it costs you nothing — reciprocity is a long-term investment.
4. Pick your fights wisely — you can't win on everything, so focus on what matters.
5. If you're new to the group, concentrate on building relationships before trying to spend them.

Chapter 15 — Practical Advice for the Standards Attorney

This chapter is the judgment layer — why the things you check matter, what goes wrong when you get them wrong, and how the practical mechanics of advising on a standards engagement actually play out. Much of the conceptual framework has been laid out in earlier chapters. This chapter applies it. For the operational prompt list — the lifecycle-organized checklist you can pull up before a particular engagement — see Appendix B.

15.1 Before You Engage: The Checklist

Start With Strategy

The first question in any standards engagement isn't legal. It's strategic. Why standardize at all? What does the client hope to accomplish — define a market, slow a competitor, build a coalition, signal commitment, satisfy a procurement requirement, prevent a worse outcome elsewhere? Without a clear answer, everything that follows — forum choice, IPR policy, governance posture, exclusion strategy — is guesswork.

This sounds obvious, and it gets skipped routinely. Engineers come to legal with "we want to join organization X" or "we want to start a new working group" already framed as the decision.

Your first job is to push back one level and make sure there's a coherent strategic theory underneath the request. If there isn't, the engagement will drift, and the legal questions you're being asked to answer will be the wrong questions.

Forum choice and structure matter, but they are secondary. They are the implementation of the strategy. A client that wants to set a market floor needs a different forum than a client that wants to slow a competitor's lock-in. A client that wants to legitimize an existing implementation needs a different structure than one that wants to design something new collaboratively. The forum and the IPR posture should fall out of the strategic answer, not precede it. When the strategy is clear, the legal architecture often becomes obvious. When the strategy is fuzzy, no amount of legal analysis will fix it.

Reading the Governance Documents and IPR Policy

Once the strategic theory is in place, before your client joins a standards body or working group, read the governance documents and the IPR policy. All of them. Not just the patent policy — the bylaws, the membership agreement, the working group charter, the contribution guidelines, and the forms.

This sounds obvious. It's remarkable how often it doesn't happen. A client's engineer joins a working group, starts contributing, and nobody reviews the IP terms until there's a problem. By then, commitments have been made.

When reading, the questions that matter most are about triggers, scope, and exits. What triggers the patent commitment — mere participation, or only contribution? What scope does the commitment cover — just the working group, or the whole organization? What exclusion or disclosure mechanism exists, and what are its timelines? What happens on withdrawal, and what survives it? What tiers of membership exist, and which one is your client joining? What are the voting rules, and which decisions require supermajority or unanimity?

Read the forms too — as discussed in Chapter 10, declaration and disclosure forms can modify the policy terms in ways that aren't reflected in the policy document itself.

Understanding the Scope of Patent Commitments

Map your client's patent exposure before joining. That means having a view of which patents your client holds that might read on the working group's scope, how broad the scope actually is (a narrow technical area, or a broad category that could expand under pressure), whether the commitment reaches optional portions of the spec, whether it sweeps in patents acquired after the commitment is made, and whether it extends to affiliates.

This doesn't require a formal patent search — as discussed in Chapter 7, there's generally no duty to search. But you should have a conversation with your client's patent team about the technology area and whether there are known patents that might be implicated.

Mapping the Organizational Structure

Understand the organization's structure before engaging. Who governs? Who votes? Who chairs the working group? What's the relationship between the working group and the governing board? Is there an association management company involved, and what role do they play?

Understanding the structure tells you where the power is. A working group chair who controls the agenda has more practical influence than a board member who only sees the final deliverable. An association management company that controls the website and membership database has operational leverage that isn't visible in the governance documents.

Identifying Your Representatives and Managing Their Knowledge

Know who your company is sending to the working group. Understand their level of knowledge about your patent portfolio. Make sure they understand the disclosure obligations — that disclosure is mandatory based on actual personal knowledge, not optional, as discussed in Chapter 7.

Ensure your representatives have appropriate training on the legal issues that arise in standards participation — particularly around patents and antitrust. Engineers and product managers who participate in standards bodies are operating in a room full of competitors, making commitments that have IP consequences, and subject to antitrust constraints they may not be aware of. They need to understand, at minimum: what they can and cannot discuss with competitors, how patent disclosure and exclusion obligations work, what triggers a patent commitment, and when to escalate to legal. This doesn't require making them lawyers. It requires making sure they know enough to recognize when they're in territory that needs legal input.

If your company has multiple people attending different working groups within the same organization, coordinate among them. Patent commitments may be working-group-specific, but the organizational rules may create broader implications. And if your representative changes, make sure the replacement is briefed on the commitments already made and the positions already taken.

15.2 During the Engagement

Monitoring Contributions and Their IP Implications

Once your client is participating, monitor what's being contributed — both by your client and by others. Contributions by your client carry IP implications under the policy. Contributions by others shape the spec in ways that may implicate your client's patents.

If the spec evolves in an unexpected direction — toward your client's patent portfolio, into a scope area that wasn't anticipated — that's when you need to evaluate options: raise the concern in the working group, consider an exclusion, or withdraw before the commitment locks in.

Managing the Exclusion Window

If the policy provides an exclusion mechanism (as most RF policies do), understand the timeline and process. Know when the call for patents will happen or, if it's a rolling exclusion, know when the final deliverable is expected.

When the window opens, coordinate with the patent team. Review the spec against known patents. Decide what, if anything, to exclude. If you're going to exclude, follow the procedural requirements — identify the patents, identify the spec sections, and file the exclusion within the deadline. Late exclusions are generally not accepted, and missing the window means the commitment locks in.

If you decide not to exclude — which is the case in the vast majority of engagements — document that decision. If questions arise later, you want a record showing that you evaluated the spec and made a deliberate choice.

Internal Coordination Across a Large Organization

In a large company, the standards team, the patent team, the engineering team, and the business unit may all have different views on a given engagement. The engineer wants to contribute freely. The patent team wants to protect the portfolio. The business unit wants the standard to align with its product roadmap. The legal team wants manageable risk.

Your job is to be the person who sees all the dimensions. Ensure the engineer doesn't make IP commitments the patent team hasn't approved. Ensure the patent team doesn't block engagement the business unit needs. Ensure everyone understands the constraints of the standards process — that voting rules, due process, and consensus requirements may not produce the outcome any single internal stakeholder wants.

This coordination is ongoing, not a one-time event. The spec evolves. The business context changes. The patent portfolio changes. Regular check-ins — not just when there's a crisis — are what prevent surprises.

When to Raise Concerns and How

If you have a concern about the direction of the spec, the governance process, or the behavior of another participant, raise it early and raise it through the process. Standards bodies have mechanisms for addressing concerns — comment periods, objection processes, appeals. Using them is legitimate and expected.

What doesn't work is waiting until the spec is finalized and then objecting. The time to influence the direction is during development, not after publication. And the time to flag a patent concern is during the exclusion or disclosure window, not after the commitment has locked in.

15.3 Risk Management

Patent Portfolio Implications

Standards participation has patent portfolio implications that extend beyond the specific commitment. Participating in a royalty-free working group means committing necessary claims to that standard for free — permanently, in most cases. If your client has a significant patent portfolio in the relevant technology area, the cost of that commitment may be substantial.

Conversely, not participating means not having a voice in the standard that may define the market. And if the standard becomes widely adopted, your client may need to implement it regardless — but without having had the opportunity to influence the result.

The decision to participate is a business decision informed by legal analysis, not a legal decision alone. Your job is to make sure the decision-makers understand the IP implications of participating and not participating.

Successor and Assign Provisions

If your client acquires a company — or is acquired — understand how the patent commitments transfer. Most policies extend to patents "owned or controlled, now or at any time in the future," which means acquired patents get pulled into existing commitments.

In M&A due diligence, check whether the target company has made standards commitments. Those commitments will survive the acquisition and bind the acquiring company's patents. This is an area that's frequently overlooked in deal diligence and can create unexpected exposure post-close.

Withdrawal Mechanics

Understand how to withdraw and what survives. Most policies allow withdrawal from a working group, but the commitment to specifications finalized before withdrawal is permanent. Some policies have notice periods. Some have specific procedures.

If your client's business circumstances change — a shift in business model, an acquisition, a decision to monetize patents — and withdrawal becomes necessary, do it cleanly and in compliance with the policy terms. A messy withdrawal creates governance disputes and damages relationships with other participants. A clean withdrawal, done with appropriate notice and communication, preserves relationships even if the departure is unwelcome.

Aligning Participation, Approval, and Monitoring

Different companies take very different approaches to standards participation, and you should help your client decide where they fall on the spectrum — ideally before they join, not after.

Participation approval. Some companies require formal legal review before anyone joins a working group. Others let engineers join freely and loop in legal only when there's a specific question. Neither approach is wrong, but the implications differ. A formal approval process ensures patent exposure is evaluated before the commitment is made. A lightweight process gets your people into the room faster but risks commitments that haven't been vetted.

Patent review approach. At one end of the spectrum, some companies greenfield an engagement — they join, participate, and never review their patent portfolio against the emerging spec. They accept the patent commitment as a cost of participation and move on. At the other end, some companies conduct formal patent reviews at multiple stages: when the working group is chartered, when the spec reaches a stable draft, and again before finalization. Most companies fall somewhere in between.

The right approach depends on your client's patent portfolio, the technology area, and the business model. A company with few patents in the relevant area can afford to greenfield. A company with a significant portfolio that it monetizes needs more rigorous monitoring. Whatever the approach, be deliberate about it. An unintentional greenfield — where nobody reviews the patents because nobody thought to — is different from a deliberate business decision to accept the exposure.

Ongoing monitoring. Standards engagements run for years. Your client's patent portfolio will change during that time — through new filings, acquisitions, and evolving technology. A patent that didn't read on the spec in year one may read on it in year three as the spec evolves. Periodic check-ins between the standards team and the patent team — not just at disclosure windows but as part of regular engagement management — reduce the risk of surprises.

Patent Exclusions: The Nuclear Option

Exclusions were covered in detail in Chapter 8. From a practical advice perspective, the key point is this: exclusions are the nuclear option. They are highly disruptive to the working group and to your client's relationships with other participants. They should be used only when absolutely necessary.

As discussed in Chapter 8, some organizations use a formal call for patents — a specific notification at a defined stage of the specification process that triggers an exclusion window. Others use a rolling exclusion obligation where participants must track the spec's evolution and exclude before finalization, without a specific trigger. From a management perspective, calls for patents are significantly easier for companies to handle. They produce a concrete notification that can be routed to the patent team, evaluated, and responded to within a defined timeline. Rolling exclusions require your organization to independently monitor the spec's progress and self-initiate the review — which, in a large organization where the standards team and the patent team may not be in regular contact, often means it doesn't happen until someone realizes the spec is about to be finalized. When evaluating which organization to work in, the exclusion model is worth considering as a practical management factor.

When an exclusion is filed, it signals to the community that your client has a patent it intends to assert or monetize — which may be exactly what's happening, but the signal is disruptive regardless. The working group has to evaluate whether to design around the excluded claims, accept them as a cost, or renegotiate the scope. Any of these responses consumes time and political capital.

If your client genuinely needs to exclude, do it narrowly. Identify the specific claims and the specific spec sections. Work in good faith with the working group to find alternatives. And communicate early — don't wait until the deadline to drop a surprise exclusion. The process is designed to handle exclusions, but it handles them far better when the excluding party engages constructively rather than dropping a stack of patent numbers on the table at the last moment.

Patent Disclosures: Good Faith and Practical Timing

For RAND engagements, patent disclosures should be good-faith representations of patents that the owner reasonably believes are necessary claims. Disclose what you genuinely believe reads on the spec. Don't over-disclose — dumping a large volume of marginally relevant patents into the disclosure list creates noise that doesn't serve the community and can slow the development process.

At the same time, try to disclose as early as possible. Early disclosure gives other participants the information they need to make informed decisions about the technology and the licensing landscape. It demonstrates good faith. And it protects your client from later accusations of ambush.

The practical reality is that early disclosure is often difficult. Specs evolve. A patent that seems relevant in a working draft may not read on the final spec, and a patent that seems irrelevant early on may become directly on point as the spec stabilizes. This is why many disclosures happen late in the process — not because participants are gaming the system, but because it's genuinely hard to evaluate whether a patent is a necessary claim until the spec is close to final.

The best approach is a combination: disclose early when you can, but plan for a more thorough review as the spec stabilizes. Document your process. If questions arise later about the timing of a disclosure, you want to be able to show that you made a good-faith effort throughout the engagement, not just at the last minute.

15.4 Drafting Standards-Related Agreements

When to Use Off-the-Shelf Frameworks

For most engagements, use established frameworks — JDF, OASIS, the Community Specification License. The terms are proven, the IP is well-understood, and the startup time is minimal. The legal teams at other participating companies have already reviewed these frameworks, which means less friction in getting to agreement.

The temptation to customize is strong. Resist it unless there's a genuine need. Every customization requires negotiation. Every negotiation takes time. Every bespoke term creates a unique legal instrument that every participant's legal team has to evaluate from scratch. The overhead of customization almost never justifies the theoretical improvement in terms.

When Bespoke Agreements Are Necessary

Sometimes the established frameworks genuinely don't fit. The technology has unusual patent implications. The participant group has specific governance requirements. The engagement involves jurisdictions or industries where the standard frameworks haven't been tested.

In those cases, draft with the standard frameworks as your starting point. Use the same structure, the same concepts, the same terminology. Deviate only where you must. And document why you deviated, because every participant's attorney will ask.

Companies Trying to Reshape Standards to Match Internal Processes

One dynamic that practitioners should recognize: companies routinely try to get standards organizations to adopt policies that mirror their own internal processes. How and when they conduct patent reviews. How they approve participation. What level of company oversight they require over employees' standards activities. Who has authority to sign agreements.

This is natural — people advocate for what they know, and if your internal patent review takes 90 days, you'll push for a 90-day exclusion window. If your company requires VP-level sign-off on patent commitments, you'll push for formal notification procedures that give you time to route the approval.

The problem is that every company's internal processes are different. What works for a company with a large patent team and formal review processes doesn't work for a startup with no patent team at all. What works for a company with centralized decision-making doesn't work for one with distributed authority.

The standards organization's policies need to work for everyone — or at least for the range of participants the organization wants to attract. When you find yourself advocating for a policy change that would make your client's life easier, ask whether it would make the process harder for other participants. If it would, consider whether your client can adapt its internal processes to the organization's rules rather than the reverse. Standards policies should be designed for the ecosystem, not for any single participant.

The Evolution to Modern Modular Frameworks

As discussed in Chapter 11, the history of standards agreements is a progression from bespoke, heavily negotiated bilateral agreements through form consortium agreements to modular, non-negotiable frameworks (JDF, CSL). Each generation reduced the startup time and legal overhead.

If you're advising on a new engagement, the question isn't whether to use a modern framework — it's which one. And if a client or counterpart proposes going back to bespoke drafting for something that a modern framework can handle, push back. The time and money spent negotiating a bespoke agreement is time and money not spent on the technical work. Get to the substance faster.

15.5 Launching a Standard: A Practical Checklist

The role of theater in standards — signaling, perception, and why trust matters more than legal perfection — is covered in depth in

Chapter 10. This section focuses on the practical mechanics of getting a spec launch right.

Review the Announcement Like a Legal Document

When advising on a spec launch, review the announcement with the same care you'd review the legal terms. Make sure the public positioning is consistent with the actual IP framework. If the announcement says "open standard" but the IP terms are RAND, someone will notice. If it says "developed collaboratively" but one company controls the governance, that tension will surface. "Open" means something specific to the communities you're addressing, and if your usage doesn't match their expectations, the gap will be exploited.

Coordinate Partner Alignment Early

Coordinating a spec launch across multiple partners — each with their own legal review, their own PR team, and their own messaging preferences — is one of the most logistically challenging aspects of standards work. It's also one of the most important.

The announcement is where the standard enters the public conversation. The partner logos signal the breadth of support. The language frames how the community will evaluate the IP terms. Get it right, and the standard launches with momentum. Get it wrong, and you spend months doing damage control.

Start the coordination early. Share draft language with partner legal teams before the announcement, not after. Give people time to flag concerns. A simple, clear message that everyone can support is better than a perfect message that takes so long to negotiate that the launch window passes.

Ensure Consistency Between Marketing and Legal

The most common launch failure isn't a bad IP framework — it's a gap between what the marketing says and what the legal terms actually provide. The marketing team wants to say "open." The legal terms say "RAND with reciprocity." The press release says "60 partners." The governance gives one company unilateral control. These inconsistencies will be found, and they will be used against you.

Before the launch, have someone — ideally the attorney who knows both the IP terms and the community — read the announcement side by side with the legal documents. Flag every claim in the announcement and verify it against the actual terms. This is a concrete, mechanical step that prevents the most common perception problems.

15.6 The Operational Layer

The reasoning in this chapter is the *why*. When you need the *did-you-check*, turn to **Appendix B — Standards Engagement Checklist for Counsel**. It mirrors the lifecycle of an engagement — before you join, IPR diligence, governance, the membership agreement, internal authorization, designating participants, ongoing hygiene, when things go wrong, and exiting — as a scannable prompt list you can pull up before a particular conversation. The two are designed to be used together: read this chapter once, keep the appendix at hand.

Practice Tips

1. Read the governance documents, the IPR policy, and the forms before your client joins.
2. Ensure your representatives have appropriate training on patents and antitrust.
3. Align your participation approval, patent review, and monitoring approach before engaging.
4. Exclusions are the nuclear option — use them only when absolutely necessary, and do them narrowly.
5. Disclosures should be good-faith representations — avoid over-disclosure, but disclose as early as possible.
6. Use established frameworks (JDF, OASIS, CSL) unless there's a genuine reason for bespoke drafting.

Chapter 16 — Lessons from Two Decades of Practice

This chapter is different from the rest of the book. It's not about patent policies or governance structures or voting mechanics. It's about the soft side of the practice of law — and it may be the most important chapter in the book.

The technical frameworks matter. The legal analysis matters. But the principles that determine whether you're effective — whether your advice gets followed, whether your clients trust you, whether you build the relationships that make everything else possible — are the ones in this chapter.

Some of these emerged from standards work. Most apply to legal practice and leadership generally. They're offered as one practitioner's perspective, not as universal truths. Take what's useful. Leave what isn't.

People and Relationships

Law is easy, people are hard. The legal analysis is rarely the difficult part of any engagement. The statutes, the case law, the contract provisions — these are learnable, analyzable, and usually within a range of well-understood outcomes. What's hard is the people. Understanding their motivations. Managing their expectations. Navigating their internal politics. Earning their trust. Getting them to act on your advice. If you can do the people part, the law part follows. If you can only do the law part, you'll produce technically correct advice that nobody acts on.

Empathy is a tool, not a weakness. Understanding what everyone at the table needs — their business constraints, their internal politics, their competitive position, their personal incentives — is strategic intelligence. The more you understand about others' interests, the better you can identify trades, build alliances, and anticipate objections. This applies whether you're negotiating a deal, structuring a partnership, or managing a team. Empathy isn't about being nice. It's about being informed.

Don't go for the jugular — these are repeat players. In any professional community — legal, technical, industry — you encounter the same people again and again. The relationship you build — or destroy — in one engagement carries into the next. A reputation for fairness, competence, and good faith is the most valuable asset in any long-term career. Once lost, it doesn't come back.

Don't push people back on their heels. This is the in-the-moment cousin of the jugular point, and it's worth its own line because it gets violated constantly by smart lawyers who think they're being effective. Remember that this isn't litigation. There's no judge to rule in your favor. You need to convince the other side to come to your position willingly, and people who feel attacked dig in. They stop evaluating your argument and start defending their own. There is a real difference between marking up someone's document and explaining all the reasons they're wrong — which almost never works, and works even worse when their client is in the room — and asking a question. "That's interesting that you said that, because I was thinking it would work like this — can you help me understand?" That single move turns a confrontation into a conversation. It gives the other side room to reconsider without losing face. It's also genuinely curious, which means you sometimes find out you were the one missing something. Either way, you end up closer to a position you can both live with than you would have by being right out loud.

You need friends before you need friends. Build relationships, goodwill, and trust before you have a specific need. Support others when it costs you nothing. Show up consistently. Be helpful without being asked. When the time comes that you need support, it will be there — because people remember who helped them and who didn't.

It's all about the chits. Early in my career, I wanted something from another team and I wasn't getting anywhere with them. My manager — the son of a southern tobacco farmer and the only person who wore a three-piece suit in a company where flannel was considered formalwear — took me aside and said, "Son, it's all about the chits." Holding out both hands, he said, "If you want something in this hand, you need to cash in the chits you've earned in the other hand. And you don't have any chits." Once I got past the condescending tone, I realized it was some of the best advice I've ever received. The only way to get chits is to earn them — by showing up, by helping others, by being fair, and by doing it before you need something in return.

When you're in someone's backyard, let them know you're there and give them the opportunity to scream. If you're working across organizational boundaries — different teams, different departments, different companies — you're going to be playing on someone else's turf. Let them know you're there, what you're doing, and why. Give them every opportunity to raise concerns. Most people will welcome you and help you navigate their particular terrain. Don't ever surprise anyone, especially in front of their stakeholders.

There's plenty of credit to go around. The instinct to claim credit is natural but counterproductive. Share it generously. Help others look good in front of their management and their clients. The more credit you give away, the more people want to work with you. As Truman is often quoted: "It is amazing what you can accomplish if you do not care who gets the credit." Do the work. Help others succeed. Your reputation will compound over time, even if — especially if — you're not the one claiming it.

Negotiation and Strategy

It's Model UN, not litigation. Any time you're working in a multi-party setting — whether it's a standards body, a joint venture, an industry consortium, or a regulatory negotiation — the dynamics are diplomatic, not adversarial. You're building coalitions, managing relationships, and trading concessions across multiple dimensions simultaneously. The skills that make someone effective in a courtroom — aggressive advocacy, clever argument — actively hurt in collaborative settings. There's no judge. You can't compel a result. And the person you went after today will be across the table from you tomorrow.

Play 3D chess — lay the groundwork and let others carry your position. The most effective work in any collaborative setting happens before the meeting. Informal conversations. Alignment sessions. Understanding where the support is and where it isn't. If you've done the groundwork, others will advocate your position for you — and it carries more weight coming from them than from you. If you haven't done the groundwork, you're improvising in real time, which is a much harder game.

The most powerful word is "no" — not "no, because." In any negotiation, once you explain your reasons for saying no, the other side will solve your stated problem and come back with a revised proposal that you may find harder to reject. Sometimes the strongest position is the unexplained one. "No" is a complete sentence. Use it sparingly, but when you do, don't dilute it with rationale that invites workarounds.

This extends to how you manage your time. Just because someone raises an issue doesn't mean you need to address it. I've seen too many hours wasted on non-issues because lawyers wouldn't say no — to a client who wanted to relitigate a settled point, to a counterpart who raised a theoretical concern with no practical consequence, to an internal stakeholder who wanted a memo on a question that didn't matter. Protecting your time — and your team's time — requires the discipline to evaluate whether an issue is worth engaging on before you engage on it.

Create the rules — once they're playing on your gameboard, you've already won. Spend time shaping the environment before you need it. Whether it's a governance framework, a contract template, a process workflow, or an organizational structure — if you're the one who designed it, you understand it better than anyone, and the dynamics favor your position. The people and companies that are most effective in any collaborative context are the ones that shaped the playing field before the game started.

Focus on what truly matters — let other things slide. You can't fight every battle. You can't perfect every clause. You can't win every point. Identify the two or three issues that genuinely matter and focus your energy there. Concede gracefully on the rest. Losing on something you don't care about preserves your credibility and your relationships for the fights that count.

Drafting and Legal Craft

The range of reasonable. In any negotiation — standards, M&A, licensing, partnerships — most terms fall within a range of reasonable. The range is wide, and there are real differences within it, but the truly unreasonable provisions are rarer than you'd expect. Your job isn't to find the perfect deal. It's to identify when something falls outside the range and flag it. Whether you love what's inside the range is a different question. But distinguishing "I don't prefer this" from "this is outside the range" is a core skill for any transactional attorney.

Oxygen expands to the size of the room. Give any project, team, or initiative an open-ended scope, and the work will grow beyond what anyone anticipated. This isn't malice — it's the natural dynamic of ambitious people without external constraints. The antidote is clear scoping at the outset and the discipline to re-scope when the work drifts rather than letting it expand silently.

Terms travel, context doesn't. Language from one agreement migrates to the next — copied, adapted, carried forward. But the context that created that language — the specific negotiation, the specific problem it was solving — doesn't travel with it. Future readers see the clause and have no idea why it exists. If you're drafting, document the rationale. If you're reviewing, don't assume a clause is meaningless just because no one can explain it — but also don't be bound by it just because it's there. If a provision doesn't make sense for the current situation, fix it or remove it. Inherited language that nobody understands is a liability, not a tradition worth preserving.

Be an expert in your space — but don't be bound by precedent. Understand the common practices, the standard terms, the typical governance structures. Know why they exist and what problems they were designed to solve. Assume they were put in place by smart people with good reasons. But don't treat them as sacred. Every procedure, every term, every mechanism was designed for a specific context — and that context may not be yours. The fact that every agreement you've seen has a particular provision doesn't mean your agreement needs it. Question what you see. Understand the purpose behind the precedent, then decide whether that purpose applies to what you're doing now. The practitioners who add the most value are the ones who know the conventions well enough to know when to break them.

It's OK to be paranoid sometimes. When you're drafting any agreement, consider the worst case. What happens if a party acts in bad faith? What happens if the relationship breaks down? What happens if someone uses the process as a weapon? You don't design for bad faith as the default — that would make every agreement unusable. But you build in safety valves that address bad faith when it occurs. The agreements that survive crises are the ones that anticipated them.

The corollary is equally important: **draft for collaboration as much as for protection.** Law school teaches you to anticipate disputes. Real practice teaches you that disputes are rarer than you'd expect. Most parties engage in good faith most of the time. If you draft every agreement as if litigation is inevitable, you create friction that makes collaboration harder — and the collaboration is usually the whole point. Draft provisions that make good-faith participation easy and rewarding. Draft governance that facilitates consensus rather than just resolving conflict. Build structures that assume the best while preparing for the worst. The balance between paranoia and optimism is what makes an agreement actually work in practice.

Contracts are the starting point for negotiations when things go wrong. A mentor taught me this early in my career, and it's shaped how I approach every agreement since. Nobody reads the contract when things are going well. The contract matters when the relationship breaks down — and at that point, it's the starting point for the conversation about what happens next, not the final answer.

This means the agreement doesn't need to be perfect. Litigation around contracts is rare. The relationships aren't. Most disputes get resolved through conversation, not courtrooms — and that conversation starts with what the contract says, even if it doesn't end there. So draft with care, but don't let the pursuit of perfection delay the work or damage the relationship. An imperfect agreement that everyone signs and moves forward with is worth more than a perfect agreement that takes six months to negotiate. That said, don't be cavalier either. The provisions that matter most are the ones that govern failure: withdrawal, dispute resolution, IP survival, and termination. Get those right. Let the rest be good enough.

A lot of what we do is theater. This applies far beyond standards. In any client-facing practice, perception matters as much as substance. How you present advice, how you frame a risk, how you position a deal — these are acts of communication that shape whether your advice gets followed. Sometimes the legally optimal approach isn't the best approach. Sometimes a simpler framework that people trust matters more than a complex one that's technically superior. Understanding the theater doesn't make you cynical. It makes you effective.

Make your legal advice human-readable. Most clients don't want a five-page memo outlining legal risks. They want to know what to do. We transformed our open source license guidance from lawyerly analysis that left engineers wondering what to do into short checklists that made calls. The volume of requests we could handle went from hundreds per month to hundreds of thousands — not because the law changed, but because we made the advice actionable. If your advice requires a law degree to understand, it's not advice. It's a research paper.

Send a demand scroll, not a nasty-gram. Lawyers default to the demand letter — "We represent Client X, here's the detailed legal analysis of everything you've done wrong, you must stop immediately, and we're prepared to sue." It's the reflex. It's also rarely the best tool. Bud Light set a better standard. When a small brewery started using the "Dilly Dilly" trademark, Bud Light didn't send a nasty-gram. They sent a costumed town crier to deliver a "demand scroll" on parchment, politely asking the company to stop — and invited them to the Super Bowl. Bud Light preserved its rights, got the infringement to stop, generated terrific PR, and made a friend instead of an enemy. Before you fire off the default cease-and-desist, ask whether there's a Dilly Dilly version of the same communication. You can almost always preserve the legal position while choosing a tone that leaves the relationship — and your client's reputation — better off.

I don't want your five-page memo — I want your judgment. This is what I tell outside counsel. I don't need a comprehensive survey of the legal landscape. I need your judgment call. What would you do? What's the risk? Is it worth taking? The ability to synthesize complexity into a clear recommendation — rather than presenting options and leaving the decision to the client — is what distinguishes a trusted advisor from a research service.

Leadership and Career

All ideas are bad — but you need 100 bad ones to get to one good one. This is about leadership and creating an environment where people contribute. If the bar for sharing an idea is high — if every proposal has to be polished and defensible — people hold back. They self-censor. They wait for someone else to go first. But if the starting assumption is that all ideas are bad, the bar drops to zero. Everyone contributes, because there's no risk in sharing a bad idea when that's the expectation. And if someone critiques your idea, it's fine — it was a bad idea anyway. No harm done. The magic is that with enough bad ideas on the table, patterns emerge. Fragments from one bad idea combine with fragments from another. Someone builds on something that was thrown out casually. And eventually, a good idea surfaces — one that nobody would have reached if the first person in the room hadn't been willing to say something imperfect. Good ideas are hard to come by. The way to find them is to make it safe to generate a lot of bad ones first.

You have apparent authority — don't be afraid to use it. As a subject matter expert, you often have more influence than your title suggests. People seek you out because you add unique value. They defer to your judgment because you've been right before. That apparent authority is real, even if it isn't on an org chart. Use it. Make decisions when decisions need to be made. If you wait for someone with formal authority to act, you'll wait too long.

Strive to be best in class. Not just good enough. Not just competent. Best in class. In any specialized legal field, being the person that others turn to — inside and outside your organization — is the most durable form of professional security. It takes years to build. It's maintained through continuous learning, generous engagement, and consistent delivery. There's no shortcut.

Give generously of your time and expertise. When someone asks for help — even if there's no direct benefit to you at the time — help them. Share what you know. Make the introduction. Review the document. Take the call. Not because you're keeping score, but because generosity builds the reputation and relationships that sustain a career. The people you help today become the people who think of you when an opportunity comes up, who vouch for you when your name is mentioned, who return the favor when you need it. What goes around comes around — not transactionally, but over the arc of a career.

Develop expertise — build a practice, don't wait for a mandate. Don't limit yourself to what's in your job description. If you see an emerging area that nobody is covering — and that needs covering — cover it. One of the most impressive things I've seen is a colleague who became interested in open data. It wasn't formally part of her job. She learned about the emerging topic, engaged across the company, and quickly became the legal open data expert at one of the largest companies in the world — because she decided to.

If you manage ambitious people, give them room to run. If you're a manager and someone on your team is building something new — exploring an area nobody asked them to explore, developing expertise that doesn't fit neatly into their job description — give them the space and protection to do it. Take away half their formal job and let them create the job they want. Maybe they'll fail, and that's OK. One home run can change everything for the better — for them, for the team, for the organization. I always told my team they should spend 20% of their time on experiments, new ideas, and stretch goals. Most didn't. So when you find someone who does, embrace them. Protect them. Clear the path. These are the people who will do great things if you let them — and who will leave if you don't.

Don't be confined to your wheelhouse — make your clients smarter. Lawyers often think they need to confine themselves to "legal" issues. The hard truth is that most of your clients don't care about indemnification and warranty clauses. That's a lawyer problem they expect the lawyer to handle. Your job is to make your clients smarter — whether they're lawyers, engineers, or business people. Flag trends. Flag competitive issues. Help them see the deal in a larger context. Help them understand why a particular point may be important to the other side. Clients come to lawyers when they have to. They come back when you make them better at their job. Even as a junior attorney, you've probably seen more of these deals than your clients have. Share that perspective. Help them understand the patterns, the risks, and the opportunities that come from having a broader view. You get your foot in the door because your clients have to call their lawyer. You get welcomed back and become a trusted advisor because you made them smarter.

Practice Tips

This chapter is itself a collection of practice principles. Treat every section heading as a tip.

Chapter 17 — Crystal Ball: Where Standards and Open Source Are Headed

A word of warning before we begin: this chapter is unlikely to age well. Predictions about technology, markets, and governance models have a poor track record, and there's no reason to believe these will be different. What follows is an honest assessment of where the forces seem to be pointing as of this writing — informed by two decades of pattern recognition, but subject to all the limitations of trying to see around corners.

Read it as a framework for thinking about the future, not as a forecast you should bet on.

17.1 Five Threads That Got Me Thinking

Several unrelated developments, when pulled together, paint a picture of where standards may be headed.

Code-first standards. Open source code is moving from being an implementation of a standard to being the standard itself. This has been happening for a while, but it's accelerating. The launch of the Agentic AI Foundation at the Linux Foundation is a recent example. Its work on the Model Context Protocol (MCP) should have been done as a spec developed with standards best practices. Instead, it's using a traditional open source approach — and it's one of the fastest-growing projects at the Linux Foundation. The tension isn't about legal structure or development model. It's that too many companies want a seat on the steering committee.

AI-generated code is production-ready. We've crossed a threshold. With the release of recent large language models, we've moved from "AI slop" to high-quality, production-ready code. One major side effect: maintainers are being overwhelmed by AI-generated pull requests — particularly around security — at a volume that human-centric governance wasn't designed to handle.

Machines are reading documentation, not people. In a conversation with the CEO of a company in the API space, he mentioned that traffic to his docs site is down over 90% in recent months — and he's hearing the same from across the industry. Developers are no longer reading documentation. They're asking AI to implement code to interact with the API, and the AI figures it out.

Reverse-engineering without specs. A developer wanted to control his robot vacuum with a video game controller. There was no documented API. He used AI to figure out the interface, build the code, and enable the controller. In the process, he discovered a security vulnerability that gave him access to the cameras and sensors of 25,000 vacuums globally. What does this mean for standards if an AI can figure out interoperability on its own?

From configuration files to self-healing code. This was an eye-opener. A developer created a system where instead of changing variables in configuration files to customize a setup, you tell the AI what you want — "switch my chat interface from WhatsApp to Slack" — and the AI figures out how to interoperate with Slack, writes new code, replaces the old code, and deploys automatically. And it's self-healing: if interop breaks, the AI detects it, figures out why, writes a fix, and redeploys. The code is automatically rewriting itself and adapting to changes.

17.2 Pulling the Threads Together

Traditionally, standards development was slow, and the goal of most interoperability standards was to achieve long-term stability. That stability was necessary to support large-scale implementations, encourage adoption, and avoid breaking changes.

That approach will likely remain valid for large-scale infrastructure — telecom protocols, power grid standards, automotive safety communications. The cost of getting those wrong is measured in lives and dollars. You're not going to let an AI dynamically rewrite a 5G base station's protocol stack in production.

But what about more specialized or application-level areas? What happens when the "standard" is open source code, the docs are generated by AI, read by AI, implemented by AI, deployed by AI, and fixed by AI when things break or change? If the implementation layer is fluid enough to adapt to breaking changes in near real-time, do you need stable, slow-moving standards?

The large-scale plumbing — protocols, transport mechanisms, languages — will likely continue to follow traditional models. It's the service level where things fragment. The roads are going to look pretty much the same. The vehicles on the road are going to be very different. It might be more like an automated middleware system, where AI figures out how to connect various endpoints dynamically.

17.3 Where the Lines Might Be

The question isn't whether AI changes standards. It's where the lines fall between what still needs human-governed standards and what doesn't.

Behavioral and regulatory standards — ISO 9000-type management standards, safety standards, AI governance frameworks, codes of practice — are fundamentally about human judgment, societal values, and regulatory intent. AI can help draft them, but the substance requires human deliberation. It's hard to automate "what level of risk is acceptable" or "what does fairness mean in this context." These aren't interoperability problems. They're policy problems. They're durable.

Large-scale infrastructure standards — telecom, power grids, automotive safety — need stability and the governance mechanisms that come with formal standards processes. The cost of failure is too high for dynamic adaptation. These are durable too.

Software and application-level interface standards — this is where AI hits hardest. If an AI can read your API, write an adapter, test it, deploy it, and fix it when it breaks, what exactly is the standard adding? For this category, the value may shift from *how to interoperate* to *what you can rely on when you do*. The AI can figure out the interface. It can't figure out the SLA, the patent commitment, or the liability allocation.

The more a standard is about technical interfaces, the less durable it may be in an AI world. The more it's about governance, liability, and rights, the more durable it is.

17.4 The IP Question Gets Harder

Here's something I've struggled with for a while. Conceptually, standards should be the patent-safe zone. They have patent commitments covering the entire standard, and they deliberately don't get into implementation details. Open source should be the patent-dangerous zone — the code is visible, it's all about implementation, and patent commitments are generally more limited.

But the reality is almost exactly backwards. There's significant patent litigation around standards — at least in wireless and codecs — and next to none around open source. Patent attorneys prefer looking at specs because specs describe what is being done in human-readable terms. Code tells you how, and mapping that to patent claims is painstaking work.

Does AI change that equation? Maybe, in two directions.

First, AI might make finding and proving infringement in code — whether open source or AI-generated — significantly easier. If an AI can read a codebase and map functionality to patent claims faster than a team of attorneys, the practical shield that open source has enjoyed starts to erode.

Second, AI might enable systematic design-around. If you can rapidly generate and test thousands of alternative implementations, you could systematically avoid patented claims. That changes the economics of design-around from expensive and slow to cheap and fast.

But there's a wrinkle. We typically advise engineers not to review patents because knowledge of a patent can increase the risk of a willful infringement finding and enhanced damages. The standard for willfulness has evolved — the Supreme Court's 2016 decision in *Halo Electronics v. Pulse Electronics* moved away from a rigid test toward a more flexible, conduct-based inquiry — but the practical advice remains cautious. If an AI is trained on patent data and generates code with awareness of what to avoid — is that willful infringement? Does the AI's "knowledge" get imputed to the user? The answer isn't clear, but the question is going to land on someone's desk sooner than we think.

17.5 The Governance Question

If AI removes the need for some categories of interoperability standards, where does governance go?

One possibility: the governance moves to the AI model itself. Whoever controls the dominant model controls the interface. Does it work better with its favored partners? Are we going from standards-based due process to open source benevolent dictators to an AI ghost in the machine? At least the benevolent dictator had a name and a mailing list.

Another possibility: standards come back in a different form. Not as interface specifications, but as behavioral standards for AI-generated systems. Think of building codes. A building code doesn't tell you where to put the kitchen. It tells you the load-bearing walls need to hold a certain weight and the electrical needs to meet certain specifications. Within those parameters, you can do whatever you want. The inspector doesn't care about your floor plan. They care about the parameters.

A model where AI generates whatever custom code it wants — but has to stay within defined parameters like API contracts, security baselines, data handling rules, and performance thresholds — might be more enforceable than what we have now. The AI can continuously validate compliance rather than relying on a human to read a 400-page spec and hope they got it right.

17.6 What Companies Will Do

Companies will always defend their business models. The power utility companies fought for years to keep Internet Protocol companies off electricity meters through standards. Phone companies fought to keep their equipment rental business. But few companies navigate major technology transitions well.

Standards organizations also need to examine their processes.If the pace of AI-driven development continues to accelerate, the traditional multi-year standards development cycle may be too slow for some categories of work. Organizations that can't adapt their processes risk becoming irrelevant — not because standards don't matter, but because the market moved while the committee was still debating scope.

17.7 What This Means for Practitioners

The IP frameworks discussed in this book remain foundational. RAND, royalty-free, necessary claims, exclusions, non-asserts — these concepts don't become obsolete because AI changes the landscape. The new artifacts may require new applications of these frameworks, but the principles endure.

Governance complexity will increase, not decrease. Adding AI governance concerns to existing standards and open source governance creates new layers of decision-making, new stakeholder dynamics, and new regulatory pressures.

The definitional battles are coming. What counts as "open" for AI? What constitutes a "standard" when the code is the spec? What are the patent implications of AI-generated implementations? Standards practitioners will be asked to weigh in on these questions, and the answers will have commercial and regulatory consequences.

But just because there's not a traditional project to control doesn't mean there isn't control. Someone trains the model. Someone curates the training data. Someone decides what "good interop" looks like when the model generates an adapter. The control doesn't disappear — it just gets laundered through a layer of abstraction that makes it harder to see. And harder to see means harder to govern.

The future is uncertain. The principles in this book are not. Apply them with judgment, adapt them as the landscape changes, and don't be surprised when the details turn out differently than anyone — including the author — predicted.

Practice Tips

1. The IP frameworks in this book remain foundational even as AI changes the landscape.
2. The more a standard is about technical interfaces, the less durable it may be in an AI world; the more it's about governance, liability, and rights, the more durable it is.
3. Watch for control that gets laundered through abstraction — someone trains the model, someone curates the data,

someone decides what "good interop" looks like.

Chapter 18 — Closing Thoughts: What It All Adds Up To

When I started the Coffee & Standards series inside Microsoft, the goal was modest — help a handful of attorneys and engineers understand the machinery of standards well enough to spot the issues that mattered. What came out of those sessions turned into this book, and along the way the material got broader than I expected. We covered IP policies. Governance structures. Voting mechanics. Open source licensing. Antitrust. Due process. The convergence of standards and open source. The looming impact of AI. Twenty years of maxims about law, people, and practice.

That's a lot of ground. And it's worth pausing at the end to ask what it all adds up to.

18.1 The Arc of the Practice

Standards work looks technical from the outside — IPR policies, balloting procedures, normative references, necessary claims. But underneath the technical surface, every chapter of this book is really about the same handful of questions.

Who gets to decide? Governance design — member classes, voting thresholds, working group charters, appeals procedures — is fundamentally about how authority gets distributed among participants with different interests. Get this right and the process is productive. Get it wrong and you get paralysis, dominance, or exit.

Who bears the cost? Patent policy — RAND, royalty-free, non-assert, exclusion — is the economic architecture of a standard. It determines who can implement, on what terms, and how the value flows between innovators and adopters. The IP framework isn't a legal afterthought. It's the commercial engine.

Who bears the risk? Antitrust rules, due process criteria, disclosure obligations, and confidentiality provisions exist because standards activity brings competitors together in a room. The rules exist to ensure that collaboration stays on the right side of the line between pro-competitive coordination and anticompetitive behavior.

Who gets the legitimacy? The same due process features that protect against antitrust challenge also unlock government recognition — A-119, Regulation 1025/2012, the TBT principles. The process that earns legitimacy in one forum tends to earn it in all of them.

Every topic in the book reduces to some combination of these four questions. If you recognize them when they show up, you'll see the structure underneath any engagement.

18.2 The Themes That Kept Recurring

Rereading the chapters, a few themes run through almost every topic. They're worth naming.

Process is substance. How a decision gets made shapes what gets decided. The governance design isn't neutral plumbing — it encodes who has leverage, what tradeoffs are easy, and what objections get taken seriously. Practitioners who treat governance as an afterthought miss most of what's happening.

Standards commoditize; they don't revolutionize. Standards take a technology that's already largely settled and turn it into plumbing. The revolution happens somewhere else — in products, in research, in open source projects that nobody standardized. By the time a standard exists, the disruptive part is usually over. Understanding this keeps expectations calibrated.

The IP framework is the heart of it. If you only had time to understand one thing about a standards engagement, it would be the patent policy. Everything else — governance, process, voting — can be fixed. A broken IP policy, especially one that's been committed to over years of participation, is nearly impossible to unwind.

Convergence is the direction of travel. Standards and open source used to be separate worlds with separate frameworks, separate communities, and separate legal analyses. They aren't anymore. A practitioner in either world who doesn't understand the other is operating at a disadvantage — and the gap will only grow.

Optics and legitimacy matter. A standard developed through a process that looks fair carries weight that a substantively identical standard developed behind closed doors does not. This isn't cynicism; it's how the system actually works. Whether you're defending against an antitrust challenge, seeking government adoption, or trying to attract implementers, the perceived legitimacy of the process is part of the output.

Judgment beats formalism. Rules help, but every rule has edge cases. The practitioners who add value are the ones who understand the purpose behind each rule well enough to know when to apply it, when to adapt it, and when to recognize that the situation is outside the rule's contemplation.

18.3 What Standards Work Teaches You

There's a question I get asked a lot by lawyers earlier in their careers: is standards work a good place to spend time? My answer is yes — and not primarily for the reasons people expect.

Standards work will teach you IP. It will teach you antitrust. It will teach you governance design, contract drafting, and organizational strategy. Those are real substantive skills, and there are worse places to learn them.

But the durable lessons from standards work aren't the doctrinal ones. They're the ones about how institutions actually function, how people behave in groups, and how decisions get made in rooms full of people with different interests.

You learn that relationships are the substrate everything else runs on — that the person across the table today will be across the table again next year, and your reputation carries from one engagement to the next. You learn that diplomacy is a professional skill, not just a personality trait. You learn to distinguish what you want to argue from what you need to win, and to save your capital for the battles that matter. You learn that process and substance are the same thing seen from two angles.

You also learn humility. Standards work surrounds you with technical people who are very good at what they do, from industries you've never worked in, solving problems you don't fully understand. The lawyers who thrive in this environment are the ones who keep learning — who make their clients smarter and let their clients make them smarter in return. The lawyers who struggle are the ones who stay in their wheelhouse.

None of this is unique to standards. What's unique is that standards work concentrates all of it in a relatively small professional community with unusually long time horizons. It's the kind of environment where you can spend a career and still be learning. That's rare, and it's worth appreciating.

18.4 Coffee & Standards

The original sessions that seeded this book were called Coffee & Standards because that's literally what they were — a standing invite to grab coffee and talk through whatever was on my mind that week. No slides. No required reading. Just a conversation about the practice.

I've tried to keep that feel in the book. Some chapters are more technical than others. Some are more personal. The goal throughout was not to produce a treatise but to produce something that reads the way I'd explain this material to a colleague at the coffee machine — with the rough edges left in, the analogies deployed freely, and the judgment calls flagged as judgment calls rather than dressed up as settled law.

If you've read this far, a few parting thoughts:

Treat this book as one practitioner's perspective, not the last word. The range of reasonable in standards work is wide. Other practitioners will disagree with positions taken here, and some of those disagreements will be well-founded. Read, adapt, and form your own view.

Don't confuse the map with the territory. This book describes a system of IP policies, governance structures, and procedural rules. The system is real, but it's not self-executing. The actual work happens through conversations, drafts, meetings, and relationships. The frameworks in these chapters are tools for navigating the actual work. They are not substitutes for it.

The people make it worth doing. Over two decades I've been fortunate to work with engineers, attorneys, policy staff, and business leaders across dozens of organizations and many countries. The technical problems are interesting. The institutional puzzles are fascinating. But the people — their generosity, their humor, their willingness to engage across differences of company and country and interest — are the reason this practice is something I'd happily do all over again.

If there's a central claim in this book, it's that standards work rewards practitioners who combine substantive rigor with human judgment, and who understand that the two are not in tension but in harmony. The frameworks are the scaffolding. The people are the building. Both matter.

Thanks for reading. Happy to chat.

— David

Practice Tips

1. Judgment beats formalism — the rules in this book are tools, not scripts.
2. Lead with the bottom line and explain tradeoffs; let decision-makers decide.
3. Optics matter — apply the reasonable-outsider test to every commitment, announcement, and process choice.
4. Consistency over purity — an imperfect position applied consistently beats an elegant one applied unevenly.
5. Play the long game — these are repeat players, and your credibility compounds across engagements.

Appendix A — Glossary of

Standards Terminology

Term	Definition
A-119	OMB Circular A-119. The US federal policy document that operationalizes NTTAA, directing agencies to use voluntary consensus standards in procurement and regulation, and defining the criteria a standards body must meet to qualify.
AACS	Advanced Access Content System. Content protection technology used in Blu-ray and HD DVD, employing per-device encryption keys with revocation capability.
ANSI	American National Standards Institute. The US body that accredits standards development organizations. Does not write standards itself.
ANSI Essential Requirements	ANSI's articulation of due process criteria for standards development: openness, lack of dominance, balance, consensus, right to appeal, notification, and consideration of views and objections. Closely parallels the A-119 criteria.
Big-S Standard	A standard from a formal international standards body (ISO, IEC, ITU). Carries treaty-level weight in trade and government procurement.

BRM	Ballot Resolution Meeting. An ISO/IEC process step where comments on a failed ballot are addressed before a re-vote.
Call for Patents	A formal notification issued by a standards body at a defined stage of specification development, triggering a window for participants to disclose or exclude patents.
CEN	European Committee for Standardization. One of three European standards organizations recognized by the EU.
CENELEC	European Committee for Electrotechnical Standardization. European standards body focused on electrical engineering.
Commitment to License	A promise by a patent holder to offer a license on specified terms (e.g., RAND or royalty-free) to implementers of a standard. Distinct from an actual license — the license itself is negotiated bilaterally.
Conformance Program	A testing and certification regime that verifies whether an implementation complies with a specification. May involve self-certification or third-party testing.
Consortium	A group of companies collaborating on a standard or specification, often through a multi-party contract rather than a formal corporate entity.
Contribution Trigger	A patent policy mechanism where the commitment is tied to technology you

	actually contribute, rather than to participation in the working group. Compare with Participation Trigger.
CSL	Community Specification License. A lightweight IP framework developed by the Joint Development Foundation for collaborative specification development in Git-based workflows.
CSS	Content Scramble System. The copy protection technology used on DVDs, predecessor to AACS.
De Facto Standard	A standard that achieved dominance through market adoption rather than through a formal standards process (e.g., Win32 API).
De Jure Standard	A standard formally recognized by an official standards body through a defined process.
Defensive Termination	A provision in a patent commitment that allows the committing party to revoke the commitment against a party that asserts patents against them.
Disclosure Obligation	A requirement in RAND patent policies for participants to notify the standards body of patents they believe may be necessary claims reading on the specification.
ETSI	European Telecommunications Standards Institute. Develops telecom standards, including cellular standards as part of the 3GPP framework.

Exclusion	In a royalty-free patent policy, the mechanism by which a participant declares specific patent claims that will not be subject to the royalty-free commitment. Considered the "nuclear option."
FRAND	Fair, Reasonable, and Non-Discriminatory. The European term for RAND licensing commitments. The F does not mean "free."
Foundation	An incorporated entity (typically 501(c)(6) in the US) that hosts standards or open source work. Provides governance, operational infrastructure, and legal identity.
Harmonised Standard	A European standard developed by CEN, CENELEC, or ETSI in response to a standardisation request from the European Commission in support of EU legislation. When cited in the Official Journal, conformity with the standard creates a presumption of conformity with the underlying EU legal requirements.
IBR	Incorporation by Reference. The regulatory technique of making a privately developed standard a legal requirement by citing it in a regulation rather than reprinting its text. Widely used in US safety and technical regulation and enabled for voluntary consensus standards under A-119.
IEC	International Electrotechnical Commission. International standards body focused on

	electrical and electronic technology. Collaborates with ISO through JTC-1.
IETF	Internet Engineering Task Force. Develops Internet protocols (TCP/IP, HTTP, TLS). Operates under the Internet Society. Uses "rough consensus and running code" as its development philosophy.
INCITS	InterNational Committee for Information Technology Standards. The US national body for ISO/IEC JTC-1 work, operating under ANSI accreditation.
IPR Policy	Intellectual Property Rights Policy. The patent and copyright rules governing a standards body, including commitment type (RAND, RF), scope, exclusion mechanisms, and disclosure obligations.
ISO	International Organization for Standardization. The largest international standards body. Works with IEC on technology standards through JTC-1.
ITU-T	International Telecommunication Union – Telecommunication Standardization Sector. UN-chartered body focused on telecom standards.
JDF	Joint Development Foundation. A 501(c)(6) nonprofit organized in Washington State that provides a "consortium-in-a-box" framework for standards projects. Now part of the Linux Foundation. Uses a Delaware Series LLC

	structure (JDF Projects LLC) to provide structural insulation between projects.
JTC-1	Joint Technical Committee 1 of ISO/IEC. The primary venue for international technology standards, including MPEG, JPEG, and related work.
Necessary Claims	Patent claims that cannot be avoided when implementing a standard. Also called "essential claims." The patent claims to which RAND or royalty-free commitments apply.
Non-Assert	A commitment by a patent holder not to enforce specified patents against implementers of a specification. Functionally similar to a license for patent purposes.
Normative	The portions of a specification that define requirements an implementation must follow. Normative text is what patent commitments typically cover. Compare with Informative.
NTTAA	National Technology Transfer and Advancement Act of 1995. US statute directing federal agencies to use voluntary consensus standards developed by private-sector bodies in lieu of government-unique standards, except where inconsistent with law or otherwise impractical. Implemented through OMB Circular A-119.
Informative	Portions of a specification that provide explanation, examples, or guidance but do

	not create compliance requirements. Patent commitments generally do not extend to informative text.
OASIS	Organization for the Advancement of Structured Information Standards. Develops enterprise, security, and interoperability standards. Offers multiple IPR modes (RAND, RF on RAND Terms, RF on Limited Terms).
OWF	Open Web Foundation. Developed community-based non-assert agreements (OWFa) for standards and specification licensing, including a CLA and Final Specification Agreement.
OSP	Open Specification Promise. A patent non-assert developed by Microsoft, structured as a "promise" rather than a license.
Participation Trigger	A patent policy mechanism where joining a working group creates the patent commitment, regardless of whether you contribute technology. Compare with Contribution Trigger.
PAS	Publicly Available Specification. A process that allows a consortium-developed specification to be submitted to ISO/IEC JTC-1 for an international vote, enabling consortium specs to achieve formal international standard status.

Patent Pool	An arrangement where multiple patent holders aggregate their standard-essential patents into a single licensing program, simplifying licensing for implementers. Administered by entities like MPEG LA or Via Licensing.
Presumption of Conformity	Under the EU's New Legislative Framework, the legal effect by which an implementer that conforms to a harmonised standard cited in the Official Journal is presumed to comply with the essential requirements of the underlying EU directive or regulation.
Profile	A defined configuration of a standard for a specific use case. Profiles narrow the options in a broadly written specification to ensure interoperability within a particular context (e.g., web streaming, satellite broadcast).
RAND	Reasonable and Non-Discriminatory. A patent licensing commitment that allows royalties but constrains the patent holder to offer licenses on reasonable terms to all implementers.
RAND-RF	RAND Royalty-Free. A RAND commitment where the royalty rate is zero. Also called RAND-Z (RAND Zero).
Reciprocity	A provision in a patent commitment that conditions the commitment on the recipient also offering patent rights back. May mean

	"back to the licensor" or "back to the world," depending on the policy.
Reference Implementation	Working code that demonstrates a specification can be implemented. Traditionally used for verification, not production. Increasingly developed as open source intended for deployment.
Rolling Exclusion	An exclusion mechanism without a formal call for patents. Participants must independently track the specification and declare exclusions before finalization.
SDO	Standards Development Organization. Broad term covering organizations that develop standards, from international bodies to industry consortia.
SEP	Standard-Essential Patent. A patent containing one or more claims that are necessary to implement a standard. The patent-level equivalent of necessary claims.
Series LLC	A legal structure where a single LLC contains multiple "series," each operating as a separate legal entity with its own assets and liabilities. Used by JDF to provide structural insulation between projects.
SIG	Special Interest Group. A small, focused collaboration. The term has largely fallen out of favor.

Small-s Standard	An industry standard from a consortium or foundation (W3C, OASIS, IETF) rather than a formal international standards body.
SSO	Standards Setting Organization. Sometimes used specifically for the "big" international bodies (ISO, IEC, ITU), though usage varies.
W3C	World Wide Web Consortium. Develops web standards (HTML, CSS, Web APIs). Incorporated as its own nonprofit in 2023. Royalty-free patent policy first adopted in 2004 and updated since.
WHATWG	Web Hypertext Application Technology Working Group. Founded in 2004 by Apple, Mozilla, and Opera to develop HTML as a "living standard" after disagreements with W3C's direction.
Working Group	A subgroup within a standards organization chartered to develop a specific specification. Typically has its own scope, IPR mode, deliverables, and decision-making rules.
Zombie Commitment	A patent commitment that persists after a participant has withdrawn from a working group, covering specifications finalized before the withdrawal.

Appendix B — Standards Engagement Checklist for Counsel

This is the operational companion to Chapter 15. Chapter 15 explains the reasoning — why the things you check matter, and what goes wrong when you skip them. This appendix is the prompt list itself: scannable, lifecycle-organized, and designed to be opened before a particular engagement rather than read straight through.

It is not a substitute for judgment, and not every item applies to every engagement. Use it to make sure you've asked the questions that matter. The checklist is organized around the lifecycle of an engagement: deciding whether to join, joining, participating, and leaving.

---

1. Before You Join — Should We Engage At All?

The first question is strategic, not legal: why standardize at all? Forum choice and structure are downstream of that answer. If the strategy isn't clear, stop here.

Business Case

- ☐ Who inside the company is championing this? Is there an identified business owner who will remain engaged over the life of the standard?
- ☐ What does the company hope to accomplish — influence the technical direction, track competitive activity, enable a product, signal ecosystem commitment?
- ☐ If the engineering team isn't willing to fund the dues and participation costs, is the engagement worth doing at all?
- ☐ Is there a realistic theory of how participation translates to outcomes the company cares about?
- ☐ What's the exit story if the business rationale goes away?

The Organization

- ☐ What type of organization is it — formal SDO (ISO/IEC, ITU, ETSI), accredited national body (ANSI, INCITS), open consortium, foundation-hosted project, ad hoc group?
- ☐ Is it an incorporated entity or a contractual arrangement? If incorporated, what type (501(c)(6), LLC, foundation) and what jurisdiction?
- ☐ How long has the organization existed? What's its track record?
- ☐ Who else is a member — and what's their level of engagement?
- ☐ Are the organization's governing documents publicly available? Have you read them?

Scope of the Work

- ☐ What specific working group, project, or specification is in scope?
- ☐ Is the scope well-defined, or is it open-ended in a way that invites mission creep?
- ☐ Could membership obligations extend to working groups the company doesn't want to participate in?
- ☐ Are there upstream or downstream dependencies (normative references) that could pull in other IPR regimes?

---

2. IPR Diligence — The Core Review

Patent Policy

- ☐ What is the IPR regime — RAND, RF, non-assert, something bespoke?
- ☐ If RAND: is FRAND defined? Are defensive suspension and reciprocity provisions in the policy?
- ☐ If RF: is the royalty-free commitment contribution-triggered or participation-triggered?
- ☐ What is the definition of "necessary claims" (or "essential claims")? Is it limited to normative text? Does it cover optional features?
- ☐ Does the commitment extend to transferees of the patents?

- ☐ Is there a defensive termination provision that allows the committing party to withdraw against a party that sues them first?

Disclosure and Exclusion

- ☐ What disclosure obligations exist? Is disclosure based on actual knowledge or constructive knowledge?
- ☐ When does the "call for patents" occur, and what is the exclusion window?
- ☐ What is the exclusion mechanism — can the company carve out specific claims, or is it all-or-nothing at the working group level?
- ☐ What happens if a patent is identified after the exclusion window closes?

Copyright and Trademark

- ☐ What are the copyright rules for contributions and for the final specification?
- ☐ Who owns the specification — the organization, the contributors, or a shared license?
- ☐ Are trademark rights in the specification name, logo, or conformance mark separately licensed?
- ☐ Is there a conformance program, and what are the trademark rules around it?

---

3. Governance and Participation Rules

- ☐ What are the member classes, and what rights attach to each (voting, board eligibility, working group access)?
- ☐ How are decisions made — consensus, supermajority, simple majority, board approval? At what level of the organization?
- ☐ Are there antitrust guidelines? Are they read at meetings? Is there an antitrust officer or designated chair responsibility?
- ☐ What are the confidentiality obligations? Do they distinguish between organization-internal information and member-contributed information?
- ☐ Are meeting minutes kept? Are they made available to members?
- ☐ Is there a code of conduct, and what is the enforcement mechanism?

---

4. The Membership Agreement

- ☐ Has counsel reviewed the current version of the membership agreement, bylaws, and IPR policy? (Not just the marketing description.)
- ☐ Are there provisions that bind the company beyond the IPR policy itself — e.g., broad confidentiality, non-disparagement, publicity rights?

- ☐ What is the term, and how does renewal work? Is renewal automatic absent notice?
- ☐ What are the withdrawal provisions — notice period, effective date, consequences?
- ☐ Which obligations survive withdrawal? Patent commitments almost always survive. What else?
- ☐ Is there a governing law and forum provision? Is it acceptable?
- ☐ Are there indemnification or limitation of liability provisions running in either direction?

---

5. Internal Authorization

- ☐ Who has authority to sign the membership agreement? Does the company's signing matrix cover this type of commitment?
- ☐ Has the engagement been reviewed for conflicts with the patent portfolio — pending litigations, licensing programs, strategic assertions?
- ☐ Has the engagement been reviewed for conflicts with existing products that may read on the developing specification?
- ☐ Is antitrust clearance required (e.g., HSR analysis, or internal competition-law review)?
- ☐ Does the engagement align with the company's open source and open standards policies?

- ☐ Has the relevant IP, business, and engineering leadership signed off?

---

6. Designating Participants

- ☐ Who will represent the company in the working group? Are they an employee? Are they identified as a company representative or as an individual?
- ☐ Has the participant been briefed on the IPR policy — especially disclosure obligations and exclusion windows?
- ☐ Does the participant have authority to bind the company on technical positions? On procedural votes? On IPR commitments?
- ☐ Does the participant know what they can and cannot say about the company's patents, products, and business plans?
- ☐ Are multiple participants coordinating their positions? Who resolves disagreements?
- ☐ Is there a documented process for the participant to escalate issues to counsel?

---

7. Ongoing Hygiene

Patent and Contribution Review

- ☐ Is there a process for reviewing contributions before they are submitted — both for IP clearance and for technical fit with the company's strategy?

- ☐ Is there a process for tracking calls for patents and meeting disclosure and exclusion deadlines?
- ☐ Is there a record-keeping process for what the company contributed, when, and under what commitment?

Meetings and Communications

- ☐ Are participants reminded of antitrust guidelines before meetings?
- ☐ Are sensitive topics (pricing, market allocation, boycotts) recognized and avoided?
- ☐ Are side conversations at standards meetings treated with the same discipline as the formal sessions?
- ☐ Are written communications (email lists, chat, issue trackers) governed by the same confidentiality and antitrust rules?

Coordination

- ☐ Is there a regular check-in between counsel, the business lead, and the participants?
- ☐ Are voting positions coordinated with stakeholders before key votes?
- ☐ Are external communications about the standard (press, blog posts, conference talks) reviewed for consistency with the company's positions?

---

8. When Things Go Wrong

- ☐ Is there a procedural issue (improper voting, inadequate notice, scope expansion) that needs to be raised — and by whom?
- ☐ Is a participant behaving in a way that creates risk for the company or the organization? How is that addressed?
- ☐ Has a third party asserted that a specification infringes their patents? What is the organization's response mechanism, and what is the company's separate position?
- ☐ Has a competition authority or regulator expressed interest in the organization's work? Has counsel been notified?
- ☐ Is the scope of the work drifting in a direction the company can't support? What are the procedural options — objection, abstention, withdrawal from a working group, withdrawal from the organization?

---

9. Exiting

- ☐ Has the business decision to exit been documented, with a clear rationale?
- ☐ What notice is required, and to whom must it be given?
- ☐ Which commitments survive — patent commitments, confidentiality, copyright licenses, trademark obligations?
- ☐ Is there confidential material that must be returned or destroyed?

- ☐ Are there public-facing statements (website, press) that need to be updated?
- ☐ Has the company's continuing implementation of the standard been evaluated in light of post-withdrawal status — do patent commitments still run in the company's favor as an implementer?

---

10. A Few Cross-Cutting Reminders

- ☐ When in doubt, read the actual governing documents — not the summary, not the marketing, not the organizational chart. The answer is almost always in the text.
- ☐ Distinguish what the organization *requires* from what the organization *does in practice*. Both matter, but for different reasons.
- ☐ Distinguish what the company *must* do from what the company *should* do. The legal minimum is not always the right answer.
- ☐ When a question falls outside the range of normal — a term you haven't seen before, a mechanism that operates differently than expected, a pattern of behavior that seems off — slow down and diligence it before you sign.
- ☐ Document your decisions. The reasoning that seems obvious today will not be obvious to your successor three years from now.

Appendix C — Standards Wars: A Case Study

The principles in this book were shaped in large part by watching standards battles play out in real time. One in particular is worth preserving in detail. The HD DVD vs. Blu-ray format war showed how a single product market can splinter across physical format, codec, content protection, interactivity, and gaming-console layers — each with its own standards dynamics.

This case study is placed in an appendix rather than in the main text because the concepts it illustrates — governance, IPR, multi-party negotiation, perception management — are covered chapter by chapter in the body of the book. The case is long, narrative, and backward-looking. Read it as illustration, not as prerequisite.

C.1 The HD DVD vs. Blu-ray Format War

The format war between HD DVD and Blu-ray (2006–2008) represents one of the most significant standards competitions in consumer electronics history. It was fought simultaneously across physical formats, video codecs, content protection systems, studio alliances, interactivity platforms, and gaming consoles. Every layer of the technology stack was a separate battleground with its own dynamics and its own stakeholders. Understanding this war requires following all of them at once — which is exactly what makes it such a rich case study.

C.1.1 The Consumer Electronics Context

By the mid-2000s, high-definition televisions were becoming mainstream, but standard DVDs offered only 480 lines of resolution — far short of what modern HDTVs could display. The consumer electronics industry needed the next generation to restart the upgrade cycle: new players, new TVs, new audio equipment. Studios saw it as an opportunity to sell content people had already purchased, upgrade copy protection, and maintain the distribution windows (theatrical, airline, rental, retail) that had become major revenue drivers.

The possibility of a format war was recognized early. Industry veterans recalled the costly Betamax vs. VHS conflict. Unification talks occurred in early 2005 but failed by August due to fundamental disagreements on physical format specifications, royalty arrangements, and interactive technologies.

C.1.2 The Physical Format Layer

The formats reflected distinct approaches to balancing innovation, backward compatibility, and manufacturing cost.

HD DVD was primarily Toshiba's technology — an evolutionary improvement on existing DVD. It offered 15GB single-layer and 30GB dual-layer capacity. Critically, existing DVD production facilities could produce HD DVDs with relatively minor modifications, and players were backward compatible with regular DVDs. Toshiba leveraged this manufacturing advantage aggressively, dropping dedicated player prices below $200 by late 2007.

Blu-ray was Sony-led revolutionary technology. It offered 25GB single-layer and 50GB dual-layer — a 67% capacity advantage over HD DVD — achieved through denser pit geometry and a more tightly focused laser beam. But it required entirely new manufacturing chains, wasn't backward compatible at the disc level (though most Blu-ray players could play standard DVDs), and players launched at substantially higher prices.

The two technologies adopted fundamentally different philosophies for delivering high-definition content. Toshiba bet on advanced video compression — next-generation codecs like MPEG-4 AVC/H.264 that could deliver 1080p quality at lower bitrates, compensating for HD DVD's smaller disc capacity. Sony took the opposite approach: keep MPEG-2, the same compression technology used for DVDs, and compensate with Blu-ray's massively larger physical storage. Studios could use their existing MPEG-2 toolchains — the authoring workflows, the encoding expertise, the quality assurance processes they'd built over a decade of DVD production — and simply make the files bigger. This was a significant selling point for studios that had invested heavily in MPEG-2 infrastructure and didn't want to retool for new codecs.

Both formats continued engineering development throughout the war. Toshiba received approval for a 51GB triple-layer HD DVD specification in November 2007, while disc manufacturers reportedly demonstrated technology that could potentially extend both formats to additional layers. These developments never reached consumers due to the war's resolution, but they illustrate how technical innovation continued to shape competitive positioning even as the business dynamics were shifting.

C.1.3 The Content Ecosystem

The format war's outcome hinged as much on content availability as on technical specifications. When the formats launched in 2006, content was divided: Universal Studios was exclusively HD DVD, Sony Pictures and Disney backed Blu-ray, and Warner Bros. and Paramount initially supported both.

Both sides employed aggressive tactics to secure exclusives. Toshiba reportedly paid approximately $150 million in promotional consideration to convince Paramount and DreamWorks to drop Blu-ray support in August 2007. These exclusivity deals created a fragmented marketplace where consumers couldn't access all desired content on a single format, inhibiting adoption of either standard.

Retail channels became critical. Blockbuster announced it would stock Blu-ray exclusively after finding over 70% of high-definition rentals were Blu-ray discs. Target followed. Walmart's February 2008 decision to phase out HD DVD — as the largest DVD retailer in the United States — prompted the New York Times to publish a mock obituary for the format. Netflix began phasing out HD DVD inventory in early 2008, and UK retailer Woolworths announced it would stock only Blu-ray in its 820 stores.

The adult entertainment industry, which traditionally held significant influence in format adoption due to its sizable stake in the home video market, was actively courted by both sides. While historically influential in previous format wars — most notably VHS vs. Betamax — the adult industry's role was less decisive in this conflict compared to mainstream studio support.

C.1.4 Content Protection: AACS, BD+, and ROM Mark

Content protection significantly influenced studio support. Both formats implemented the Advanced Access Content System (AACS), which represented a major advancement over DVD's Content Scramble System (CSS). Unlike CSS, which provisioned all players of a given model with the same shared key, AACS provided each individual player with a unique set of decryption keys using a broadcast encryption scheme. This allowed content providers to revoke individual compromised players rather than entire device models — addressing a fundamental weakness that had led to CSS being cracked.

AACS also incorporated traitor tracing techniques: multiple versions of short sections of content were encrypted with different keys. By embedding varying digital watermarks in these sections, analysts could trace pirated content back to the specific compromised player. The compromised keys could then be revoked in future disc releases.

Blu-ray implemented two additional protection layers that became significant differentiators. **BD+** was a virtual machine embedded in authorized players that could examine the host environment for tampering, verify player key integrity, and transform audio and video output to prevent unauthorized viewing. BD+ was designed for renewability — content providers could include executable programs on discs that addressed security vulnerabilities even after players were in consumer hands. One analyst claimed the BD+ protection scheme would take "ten years" to crack. **ROM Mark** was an undetectable identifier embedded in pre-recorded Blu-ray media during manufacturing. Since it could only be created with licensed manufacturing equipment, the ROM Mark made bit-for-bit piracy impractical. Several studios specifically cited BD+ as their reason for supporting Blu-ray over HD DVD.

C.1.5 Managed Copy and Consumer Rights

Both formats included provisions for "Managed Copy" — a system allowing consumers to make legal copies of content while maintaining DRM protection, supporting scenarios like backups, media server storage, or scaled-down versions for portable devices.

Managed Copy became a point of contention. HD DVD included mandatory Managed Copy from the outset. Blu-ray only adopted it later, after pressure from HP. This reflected different philosophical approaches to balancing content protection with consumer convenience. HP went so far as to propose an ultimatum to the Blu-ray Disc Association: adopt Microsoft's HDi interactive platform and implement mandatory Managed Copy, or HP would support HD DVD instead. Blu-ray added Managed Copy but maintained its commitment to BD-J, and HP ultimately remained in the Blu-ray camp.

C.1.6 The Codec Layer

Given the competing compression philosophies — Toshiba betting on advanced codecs, Sony betting on bigger discs with existing MPEG-2 — the codec question was tightly coupled to the physical format battle.

Three codecs were in play. **MPEG-2** was the incumbent — the same technology used for DVDs, with a decade of toolchain investment across the content industry. Studios had mastered MPEG-2 encoding, quality assurance, and authoring workflows. Switching to a new codec meant retraining staff, rebuilding tools, and requalifying the entire production pipeline. **MPEG-4 AVC/H.264** was the next-generation international standard — dramatically more efficient than MPEG-2, capable of delivering 1080p at lower bitrates, but requiring new encoding infrastructure. **VC-1** was originally Microsoft's Windows Media Video 9, subsequently standardized through SMPTE. Microsoft had invested heavily in VC-1 and built an extensive authoring and encoding toolchain around it.

Each codec had different patent royalty structures, different toolchain maturity, and different industry backing. The question of which codec would be adopted had significant financial implications — not just in licensing fees but in the value of existing toolchain investments and the control each company would have over the encoding ecosystem.

Through extensive negotiation — driven in significant part by Warner Brothers, which had invested in VC-1 toolchains — both Blu-ray and HD DVD agreed to make all three codecs mandatory.

Every player had to implement MPEG-2, H.264, and VC-1, because there was no way to predict what codec a particular disc would use.

This resolved the codec-within-a-format conflict but at a permanent cost. Every player shipped with three codec implementations, each carrying its own patent royalties and engineering burden. The industry effectively moved to MPEG-4 AVC/H.264 as the dominant encoding format, but players still need to support all three for legacy compatibility. An old disc encoded in VC-1 still has to play. Standards accrete — once embedded, technologies persist indefinitely. Every additional requirement adds permanent cost that outlives the strategic rationale for including it.

AACS also included an Image Constraint Token (ICT) feature that could restrict analog outputs to lower resolution while allowing full 1080p only on digital outputs supporting HDCP protection. Studios implemented ICT differently, adding yet another layer of complexity for player manufacturers and consumers.

C.1.7 The Interactivity Layer

The interactive programming environments represented a significant point of differentiation — and reflected corporate alliances and strategic rivalries that extended far beyond optical media.

HD DVD's HDi platform was based on web technologies: HTML, XML, CSS, SMIL, and ECMAScript (JavaScript). Microsoft developed the HDi Interactive Format, making it relatively accessible to web developers without requiring specialized DVD authoring experience. Microsoft also released an HD DVD add-on for its Xbox 360 console, positioning it as a competitor to Sony's PlayStation 3.

Blu-ray's BD-J platform was based on Java. This decision came after considerable debate — during unification talks in 2005, Sun Microsystems announced that the Blu-ray Disc Association had chosen BD-J instead of Microsoft's HDi.

This wasn't a technical accident — it was an extension of the decade-long Microsoft vs. Sun Microsystems competition for platform control. Microsoft had historically viewed Java as a significant competitive threat. Sun's positioning of Java as platform-independent ("write once, run anywhere") directly challenged Microsoft's Windows-centric model. The conflict dated back to the late 1990s: Microsoft had created a Windows-only Java implementation (Visual J++), Sun had sued for breach of the licensing agreement, and in 2002 Sun filed a $1 billion antitrust lawsuit alleging Microsoft used its Windows monopoly to undermine Java.

By the mid-2000s, Java had established a growing presence in mobile devices (J2ME) and in interactive television through the Multimedia Home Platform (MHP), which evolved into the Globally Executable MHP (GEM) — the basis for BD-J. By 2009, GEM had reached 33 million deployments worldwide. For Microsoft, supporting HD DVD's interactivity layer was part of a broader strategy to counter Java's expansion into consumer electronics. The choice of interactivity platform for next-generation optical media was another front in that larger war —

one where the outcome would influence platform control across television, mobile, and embedded devices.

C.1.8 PlayStation 3 as Blu-ray's Trojan Horse

Perhaps the most consequential business decision was Sony's integration of Blu-ray into the PlayStation 3. By February 2008, approximately 10.5 million PS3 consoles had been sold worldwide, compared to an estimated 1 million total HD DVD players. This created a massive installed base advantage.

The strategy came at enormous cost. Sony priced the PS3 at $599 at launch — roughly double the PS2 — and sold each console at an estimated loss of more than $200 per unit, resulting in aggregate losses estimated at about $3 billion. But it worked. Blu-ray titles outsold HD DVD two-to-one in the US and by three or four to one in Europe.

Microsoft took the opposite approach: an optional HD DVD add-on for the Xbox 360. This kept the base console price low but resulted in far lower HD DVD adoption. And since the Xbox 360 didn't use HD DVD for games (unlike the PS3, which used Blu-ray discs for gaming), the add-on had limited appeal.

C.1.9 How the War Ended

Warner Bros.' January 4, 2008 announcement that it would drop HD DVD support exclusively was the decisive moment. As the largest home video distributor, Warner's defection triggered a cascade: Netflix, Best Buy, and Walmart all announced they would phase out HD DVD within weeks. Toshiba announced on

February 19, 2008 that it would cease manufacturing HD DVD players — less than two years after both formats launched.

C.1.10 What It Teaches Us

The format war offers several lessons specific to the HD DVD vs. Blu-ray dynamic:

Technical superiority doesn't guarantee success. HD DVD's lower manufacturing costs, backward compatibility, and consumer-friendly pricing kept it competitive for nearly two years despite Blu-ray's technical advantages.

Content ecosystem control can outweigh everything else. Warner Bros.' decision to go exclusive was the war's decisive turning point. One company's content catalog decision overrode years of technical development and billions in investment.

Strategic platform integration changes the math. Sony's PlayStation 3 strategy — subsidizing billions in hardware losses to build an installed base — fundamentally altered the competitive landscape.

Security features influence institutional support. Blu-ray's additional content protection layers directly secured studio backing from risk-averse content owners.

Format wars reflect broader industry conflicts. The Java vs. Microsoft interactivity battle, the patent royalty dynamics around codecs, the content protection governance through DVD-CCA — each layer of the format war was a proxy for a larger competitive struggle.

Practice Tips

1. Map every layer of the stack when entering a format war — physical, protection, codec, interactivity. A win at one layer can be undone by a loss at another.
2. Track studio, platform, and distribution alignments as carefully as the technical specs — market share follows content availability, not spec quality.
3. Remember that a subsidized anchor platform (PS3 in Blu-ray's case) can decide an otherwise evenly matched fight.
4. Be skeptical of winning the format when the underlying market is shrinking — Blu-ray won physical media just as streaming was displacing it.
5. Every layer locked into the standard persists long after its strategic rationale has faded.

About the Author

David Rudin is an Assistant General Counsel at Microsoft, where he has spent over twenty years focused on intellectual property, standards, and open source. He led Microsoft's Open Source and Standards Legal Team, working in close partnership with the company's Open Source Program Office to help Microsoft engage in open source at scale.

David has built several of the legal tools discussed in this book. He founded the **Joint Development Foundation (JDF)**, pioneering the use of Series LLC structures to create a "consortium-in-a-box" that lets new standards organizations launch in days rather than months. JDF now hosts over thirty projects — including the Alliance for Open Media, Overture Maps, and C2PA — and is part of the Linux Foundation. He led the development of the **Open Web Foundation agreements**, which brought the low friction of open source-style licensing to standards development. And he created the **Community Specification License (CSL)**, designed to bring standards development best practices into the Git-based workflows developers use daily.

As legal chair for the **Alliance for Open Media**, David built the legal infrastructure supporting the AV1 video compression standard, working with attorneys from across the technology ecosystem. AV1 now powers video across hundreds of millions of devices for platforms including Netflix, YouTube, and Teams.

This work was recognized by the Academy of Television Arts & Sciences with a Technology & Engineering Emmy® Award.

He holds a JD and a Master of Science in Management Information Systems from Boston University and a BA from the University of California, Santa Barbara. He lives in Redmond, Washington.

COFFEE & STANDARDS

"

*The judgment layer — the tradeoffs,
the patterns, and the practical knowledge
that comes from doing the work long enough
to have seen most of the problems more than once.*

Technical standards shape how the world's technology connects — from the web protocols that power the internet to the video codecs that stream your content. But the agreements behind those standards are rarely understood, even by the attorneys and engineers who work with them every day.

This is the book that didn't exist. Written by a practitioner who has spent over two decades building the legal frameworks used by organizations like the Alliance for Open Media, the Joint Development Foundation, and the Open Web Foundation, it is the operational handbook for advising on a standards engagement from first contact to final publication.

READ IT AND YOU WILL KNOW HOW TO

- Read an IPR policy and spot what's outside the range of reasonable
- Tell when a royalty-free commitment actually delivers royalty-free
- Get scope, necessary claims, and exclusions right
- Navigate multi-party negotiation among repeat players
- Recognize the parts of the work that are theater

No treatise. No academic theory. Just the judgment layer.

Written for standards and IP counsel — and for the engineers, governance professionals, and business leaders at the table with them.

DAVID RUDIN

STANDARDS & OPEN SOURCE COUNSEL

FIRST EDITION · STANDARDS · IP · GOVERNANCE