

Stay connected all on
America's largest and most reliable
3G network.

verizon wireless
TOUGHBOOK



Harvard Business Review

Which Kind of Collaboration Is Right for You?

by Gary P. Pisano and Roberto Verganti

The new leaders in innovation will be those who figure out the best way to leverage a network of outsiders.

In an era when great ideas can sprout from any corner of the world and IT has dramatically reduced the cost of accessing them, it's now conventional wisdom that virtually no company should innovate on its own. The good news is that potential partners and ways to collaborate with them have both expanded enormously in number. The bad news is that greater choice has made the perennial management challenge of selecting the best options much more difficult. Should you open up and share your intellectual property with the community? Should you nurture collaborative relationships with a few carefully selected partners? Should you harness the "wisdom of crowds"? The fervor around open models of collaboration such as crowdsourcing notwithstanding, there is no best approach to leveraging the power of outsiders. Different modes of collaboration involve different strategic trade-offs. Companies that choose the wrong mode risk falling behind in the relentless race to develop new technologies, designs, products, and services.


All too often firms jump into relationships without considering their structure and organizing principles—what we call the *collaborative architecture*. To help senior managers make better decisions about the kinds of collaboration their companies adopt, we have developed a relatively simple framework. The product of our 20 years of research and consulting in this area, it focuses on two basic questions: Given your strategy, how open or closed should your firm's network of collaborators be? And who should decide which problems the network will tackle and which solutions will be adopted?

Collaboration networks differ significantly in the degree to which membership is open to anyone who wants to join. In totally open collaboration, or crowdsourcing, everyone (suppliers, customers, designers, research institutions, inventors, students, hobbyists, and even competitors) can participate. A sponsor makes a problem public and then essentially seeks support from an unlimited number of problem solvers, who may contribute if *they* believe they have capabilities and assets to offer. Open-source software projects such as Linux, Apache, and Mozilla are examples of these networks. Closed networks, in contrast, are like private clubs. Here, you tackle the problem with one or more parties that *you* select because you deem them to have capabilities and assets crucial to the sought-after innovation.

Collaboration networks also differ fundamentally in their form of governance. In some the power to decide which problems are most important, how they'll be solved, what constitutes an acceptable solution, and which solutions should be implemented is completely vested in one firm in the network: the "kingpin." Such networks are hierarchical. Other networks are flat: The players are equal partners in the process and share the power to decide key issues.

Discussions of collaborative innovation in both academic journals and the popular media often wrongly link "openness" only with "flatness"—and even suggest that open, flat approaches are always superior. The notion is deeply flawed, however.

As the exhibit "The Four Ways to Collaborate" shows, there are four basic modes of collaboration: a closed and hierarchical network (an *elite circle*), an open and hierarchical network (an *innovation mall*), an open and flat network (an *innovation community*), and a closed and flat network (a *consortium*).

 The Four Ways to Collaborate (Located at the end of this article)

When figuring out which mode is most appropriate for a given innovation initiative, a firm should consider the trade-offs of each, weighing the modes' advantages against the associated challenges and assessing the organizational capabilities, structure, and assets required to manage those challenges. (See the exhibit "How to Choose the Best Mode of Collaboration.") Its executives should then choose the mode that best suits the firm's strategy.

 How to Choose the Best Mode of Collaboration (Located at the end of this article)

Open or Closed Network?

The costs of searching for, screening, and selecting contributors grow as the network becomes larger and can become prohibitive. So understanding when you need a small or a large number of problem solvers is crucial. Closed modes, obviously, tend to be much smaller than open ones.

When you use a closed mode, you are making two implicit bets: that you have identified the knowledge domain from which the best solution to your problem will come, and that you can pick the right collaborators in that field. Alessi, an Italian company famous for the postmodern design of its home products, bet that postmodern architecture would be a fruitful domain for generating interesting product ideas and that it could find the best people in that field to work with. It invited 200-plus collaborators from that domain to propose product designs. If you don't know where to look for solutions or who the key players are (and have no way to find out), a closed mode like Alessi's elite circle is a dangerous shot in the dark.

The big advantage of an open network is its potential to attract an extremely large number of problem solvers and, consequently, a vast number of ideas. You do *not* need to identify either the best knowledge domains or the most appropriate experts in those domains. It's like throwing an open house party: You just make it known you are having a party and provide the right inducements, and (you hope) the right people will show up.

With open participation, you don't need to know your contributors. Indeed, the fact that you *don't* know them can be particularly valuable; interesting innovative solutions can come from people or organizations you might never have imagined had something to contribute. That is the concept behind Threadless.com, a largely online retailer of T-shirts, whose designs come from the masses. By operating an innovation mall where 600,000 members submit proposals for about 800 new designs weekly, Threadless gets a steady flow of unusual and singular ideas. (Mall members and visitors to the website vote on the designs, but the Threadless staff makes the final decision on which ones to produce and rewards their creators.)

Open modes, however, have their disadvantages. Notably, they're not as effective as closed approaches in identifying and attracting the best players. That's because as the number of participants increases, the likelihood that a participant's solution will be selected (especially for an ambiguous problem) decreases. The best parties, therefore, prefer to participate in closed relationships. Open modes work best when the spread between the ideal solution and the average solution is not big and the consequences of missing out on a much better solution from an elite player are small.

Open modes are effective only under certain conditions. First, it must be possible to evaluate proposed solutions at a low cost. Sometimes the screening process is extremely cheap and fast. (For instance, it might be easy to assess whether a particular module of software code works or has bugs.) In other cases, though, the only way to find out whether an idea is worth pursuing is through expensive and time-consuming experiments, and you'll want to consider fewer (but better) ideas. The only way to do that is to invite contributions from the problem solvers that you think will have the best chance of providing good ideas. That is, to opt for a closed mode.

Consider the following simple but scary example. You have a serious illness, and you want to find the best possible treatment. Employing an open mode, you post your problem on the internet, ask for advice, and get 50 ideas that look interesting. But immediately, you face two issues. The first is what statisticians refer to as a sample selection problem: Are these the best 50 ideas out there? Maybe the most knowledgeable doctors are so busy treating patients that they don't participate in these forums, and only the doctors who have time on their hands (a bad sign for sure!) responded. The second issue is that you have to invest a lot of time and resources to evaluate the 50 ideas (visiting doctors and so on). Even worse, you may have only one shot at getting the right treatment. (Are you really going to "try out" more than one surgery?) That is why when confronted with a medical problem, we might do some research to identify elite specialists, pick one, and then seek a second opinion from one or two others.

Alessi is in a similar boat. Given the large population of designers, it could easily launch an open design competition for, say, a corkscrew on its website. With its high standing in the world of design, the firm would probably attract many proposals. However, it is not posing technical problems that have one or a few optimal solutions that can be clearly defined, thereby allowing contributors to screen many of their ideas themselves. Alessi is looking for concepts whose value is based on intangible properties such as aesthetics and emotional and symbolic content. Since there is no clear right or wrong answer, Alessi could receive thousands of proposals, creating a massive evaluation burden for the company. And because the company's strategy is to offer products with radical designs that anticipate market needs, its offerings often initially confuse consumers. Therefore it can't shift the evaluation burden to customers by asking them which designs they prefer, as Threadless does. That's why Alessi has to ensure that it will receive a few good ideas from a relative handful of contributors.

Another requirement of open modes is that participating in them must be easy. This is possible when a problem can be partitioned into small, well-defined chunks that players can work on autonomously at a fairly low cost. Someone creating a potential decoration for a Threadless T-shirt doesn't need sophisticated design infrastructure or knowledge of how the company will knit yarns or tailor shirts. The inherently modular structure of the Linux open-source community allows software developers to create code for new features without touching other parts of the application, which has more than four million lines of code. Over the past decade, such open collaboration has been made easier by information platforms that allow participants to make contributions, share work, and observe the solutions of others.

Of course, not all problems can be partitioned into small, discrete chunks. For example, the development of radically new product concepts or product architectures is an integral task that has to be embraced in its entirety. In such cases, closed modes that provide an environment where collaborators can closely interact must be employed. This is what led IBM to invite a handful of selected partners (including Siemens, Samsung, Freescale, Infineon, and STMicroelectronics) to join its Microelectronics Joint Development Alliance consortia for developing semiconductor technologies such as memory, silicon-on-insulator components, and chip-manufacturing processes.

Flat or Hierarchical Governance?

As discussed earlier, the chief distinction between a hierarchical and a flat form of governance is who gets to define the problem and choose the solution. In the hierarchical form, a specific organization has this authority, which provides it with the advantage of being able to control the direction of the innovation efforts and capture more of the innovation's value. In the flat form, these decisions are either decentralized or made jointly by some or all collaborators; the advantage here is the ability to share with others the costs, risks, and technical challenges of innovating.

Hierarchical governance is desirable when your organization has the capabilities and knowledge needed to define the problem and evaluate proposed solutions. Consider companies that post scientific problems on the innovation mall InnoCentive.com. The problems are generally smaller pieces of the sponsors' much larger R&D programs. These kingpins have a clear understanding of the relevant technologies and markets (user needs and functional requirements) and can define the system configuration and coordinate the work of various collaborators.

Conversely, flat modes work well when no single organization has the necessary breadth of perspective or capabilities. Look again at open-source software projects. These often develop very specific modules of code to address problems that users have encountered (a bug in an existing piece of code or the need for a specific hardware driver). In this case, the users are best positioned to devise and test solutions because they're closest to the problem. Indeed, they usually have discovered the problem in the first place. Or take IBM's microelectronics consortia. Since semiconductor companies other than IBM possessed critical knowledge, skills, and assets needed for microprocessor design, a hierarchical structure would have made no sense.

Flat modes are also appropriate when collaborators all have a vested interest in how a particular problem is solved and will participate only if they get some say in the decisions. For example, all the members of the IBM consortia formed over the years have expected to use in their own factories and product lines the technologies they develop collaboratively. For this reason, IBM and its partners chose a governance structure that provided each a strong voice in how the technology is developed.

Designing incentives—both financial and nonfinancial—that attract external collaborators is crucial with any of the four modes of collaboration. Nonfinancial rewards like high visibility in the job market, an enhanced reputation among a peer group, the psychological fulfillment of pursuing a strong interest, and the chance to use solutions in one's own business can replace or complement monetary rewards. There are no hard rules about which incentives work best with particular forms of collaboration. Although people

often associate psychological fulfillment with innovation communities, it can be a powerful incentive in the other modes as well. For example, Alessi not only shares royalties from sales with the designers in its elite circle but also includes their names in product marketing and offers them a high degree of freedom in the design process.

A Matter of Strategy

Choosing a collaboration mode involves more than understanding the trade-offs. A firm must take into account its strategy for building and capturing value. And as the strategy evolves, the right mode of collaboration might change, too.

Consider the approach that Apple used in developing software for the iPhone and how it changed over time. A key part of Apple's business strategy (across all its products) has been to maintain the integrity of its systems. Indeed, one of the joys (and thus differentiators) of an Apple product is that everything—the machine's hardware, software, and peripherals—seems to work together so seamlessly. Historically, this kept Apple more oriented toward closed modes, where it could better control the components that influenced the user's experience. The company took that approach in developing the first generations of the iPhone as well and relied on elite circles to develop early applications for it.

However, once the iPhone was established, Apple faced the challenge of adding software functionality and applications that would fuel more growth. Our framework helps map out the various options Apple had. It could define the applications it thought would be useful (for example, a way to synchronize the iPhone with various mobile banking systems) and then engage the best software designers to develop them (the elite circle mode again). It could partition the development of particular applications into simple chunks and then go to a bazaar like TopCoder.com and tap hordes of software developers to write code for each chunk (the innovation mall mode). It could release a development package to third-party developers and let them define and create applications that would be useful (the innovation community mode). Or Apple could work jointly with a firm like Intuit to create mobile banking software (the consortium mode). Each of these modes could certainly generate new applications, but each would have a very different impact on the iPhone platform.

To stick with the elite circle mode, Apple needed to feel confident that it knew which applications customers would want and could identify the best partners for creating them. Given the huge variety of potential applications, Apple realized that there was no way that, either alone or with a small group of collaborators, it could anticipate all the applications that an iPhone owner might find useful or just fun. So it opted to encourage a thousand flowers to bloom and allow the market to decide which ones should be picked. This reasoning ruled out the elite circle, the consortium, and the innovation mall. Accordingly, Apple introduced a kit in March 2008 that allows a community of third-party developers to create applications based on the iPhone OS platform and provide them to users directly through the iPhone. (If an application is not free, the developers keep 70% of its revenues and Apple gets 30%.)

The rollout of mobile phones using Android, Google's operating system, could prompt Apple to adopt a two-part collaboration strategy. Since Android is open-source software, it may attract an even larger community of developers than the iPhone. So Apple might decide to supplement the applications developed by third parties with proprietary hardware features conceived by its own staff and created with the help of elite circles of hardware manufacturers. That illustrates another important point: Companies can use a combination of collaboration modes simultaneously to support their strategies.

IBM's successful use of *both* an innovation community and consortia to support the strategy of its server and mainframe computer businesses is an excellent real example. IBM's strategy is to compete on the basis of hardware differentiation and service. Toward that end, the company has striven to commoditize operating systems by embracing Linux and participating actively in the open-source community—one of the first major computer makers to do so. But to continue to differentiate its hardware, IBM needs to stay on the leading edge of microprocessor technology. Given the increasing scale required to keep up with the likes of Intel, IBM turned to its consortia of semiconductor companies, which have shared development costs. This combination of innovation approaches has allowed IBM to gain market share in an intensely competitive and dynamic market.

As IBM illustrates, a key component of strategy is exploiting a firm's unique assets and capabilities. In choosing one or more collaborative modes, a firm's senior managers therefore must ask: Which of our unique assets and capabilities are we trying to enhance the value of? And what's the best way to enhance it?

A firm's collaboration capability itself can be exploited for profit. InnoCentive.com, for example, is a spin-

off of an innovation mall developed by Eli Lilly for internal purposes. Alessi is now leveraging the value of its connections with more than 200 designers by assisting companies in other businesses with product design. Alessi helps them identify the designers (usually from its own network) who can best address their specific needs. In return, Alessi receives royalties from sales of the resulting products—which now account for almost 30% of its revenues.

A New Source of Advantage

As with any strategic variable, collaborative approaches to innovation offer an array of choices and complex trade-offs. As the examples in this article suggest, each approach can be highly effective under the right conditions. Senior managers need to be wary of the notion that one type of collaboration is superior to others. Open is not always better than closed, and flat is not always better than hierarchical.

Developing an effective approach to collaboration starts with a solid understanding of your company's strategy. What is the business problem you want innovation to solve? Are you (like Alessi) trying to create a distinctive product that breaks boundaries? Are you (like IBM) trying to keep up with larger rivals (like Intel and Taiwan Semiconductor) in an intense technology race? Or are you (like Apple today) looking to broadly expand the applications of your product?

Companies must also ask what unique capabilities they bring to the collaborative process. Firms with deep relationships in a space, for example, are much better positioned to exploit an elite circle mode than a newcomer is.

It's not surprising, then, that differences in strategy and capabilities can lead to different kinds of collaboration networks competing against one another in the same industry. Thus, the task of senior leadership in innovation has broadened and become truly strategic. It is no longer just a matter of hiring the most talented and creative people or establishing the right internal environment for innovation. The new leaders in innovation will be those who can understand how to design collaboration networks and how to tap their potential.

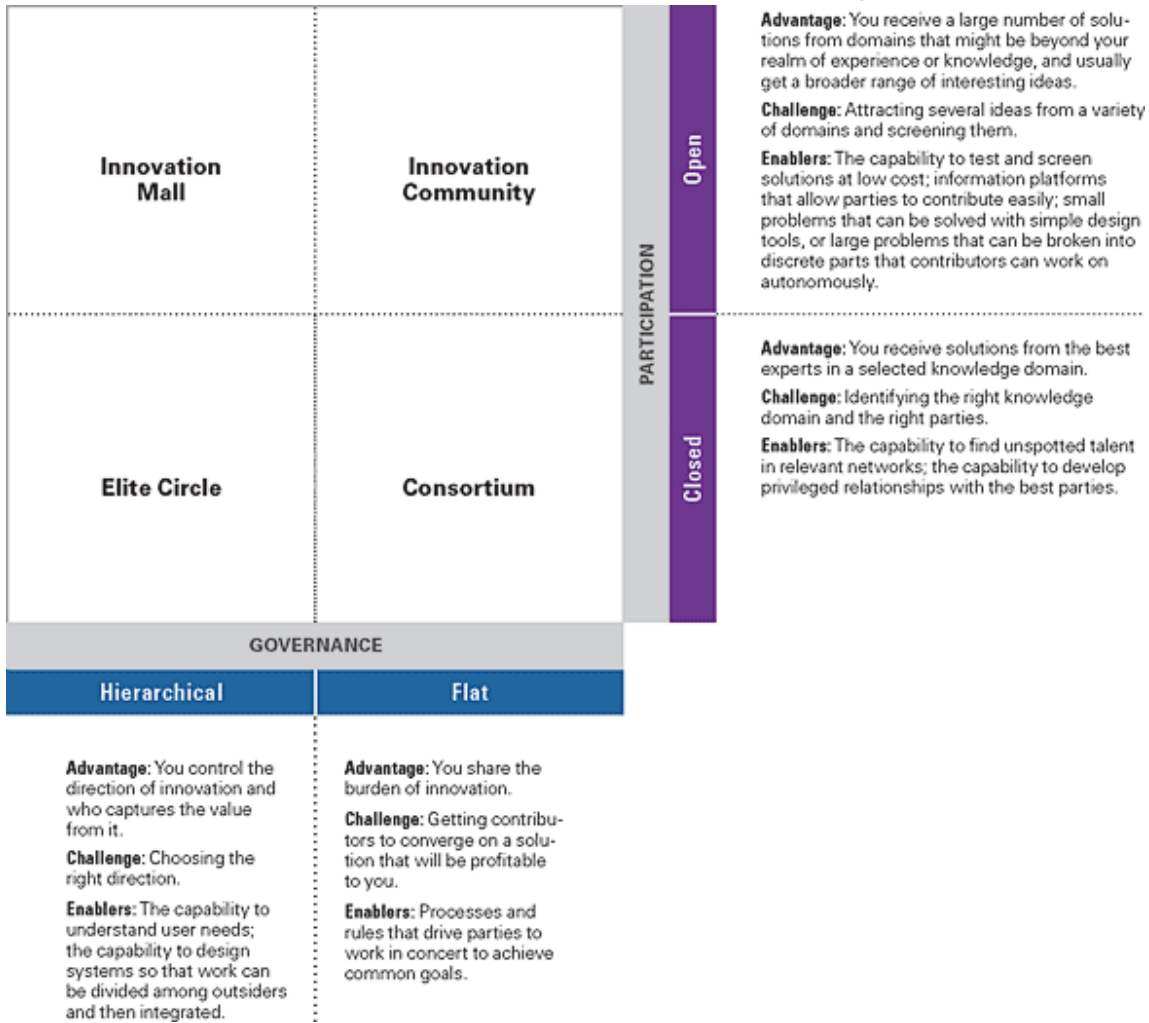
The Four Ways to Collaborate

There are two basic issues that executives should consider when deciding how to collaborate on a given innovation project: Should membership in a network be open or closed? And, should the network's governance structure for selecting problems and solutions be flat or hierarchical? This framework reveals four basic modes of collaboration.

| | | |
|--|--|----------------------------------|
| <p>Innovation Mall</p> <p>A place where a company can post a problem, anyone can propose solutions, and the company chooses the solutions it likes best</p> <p><i>Example: InnoCentive.com website, where companies can post scientific problems</i></p> | <p>Innovation Community</p> <p>A network where anybody can propose problems, offer solutions, and decide which solutions to use</p> <p><i>Example: Linux open-source software community</i></p> | <p>PARTICIPATION</p> <p>Open</p> |
| <p>Elite Circle</p> <p>A select group of participants chosen by a company that also defines the problem and picks the solutions</p> <p><i>Example: Alessi's handpicked group of 200-plus design experts, who develop new concepts for home products</i></p> | <p>Consortium</p> <p>A private group of participants that jointly select problems, decide how to conduct work, and choose solutions</p> <p><i>Example: IBM's partnerships with select companies to jointly develop semiconductor technologies</i></p> | |
| <p>GOVERNANCE</p> | | |
| <p>Hierarchical</p> | <p>Flat</p> | |

How to Choose the Best Mode of Collaboration

When selecting a mode of collaborative innovation, executives need to consider the distinct strategic trade-offs of each mode. Below are some important advantages and challenges of the different approaches to collaboration, and examples of capabilities, assets, processes, and kinds of problems that make each easier to carry out.



Copyright © 2008 Harvard Business School Publishing Corporation. All rights reserved.

Gary P. Pisano (gpisano@hbs.edu) is the Harry E. Figgie, Jr., Professor of Business Administration at Harvard Business School in Boston. **Roberto Verganti** (roberto.verganti@polimi.it) is a professor of innovation management at the Politecnico di Milano, Italy, and the author of *Design-Driven Innovation: Changing the Rules of Competition by Radically Innovating What Things Mean*, to be published by Harvard Business Press in spring 2009.