Hogan Lovells

Joint development and collaboration agreements

Primer on key contract terms

When two or more organizations want to work together to develop or enhance their products, combine or integrate their technologies, or jointly commercialize a new product, they have many choices for documenting their relationship. If the collaboration is simple and development work minimal, the parties might use standard licensing agreements and purchase orders. If the parties contemplate creating an ongoing business and making substantial investments, creation of a separate joint venture entity may be the best path. In many cases, however, a joint development or collaboration agreement provides the right framework—establishing a set of rules tailored to the relationship without the overhead and complexity of a separate joint venture. For convenience, this primer refers to joint development and collaboration agreements as "JDAs."

This primer considers key contract points that recur frequently in JDAs and is designed to provide points for consideration and checklists of items for the attorney to consider in preparing and negotiating a JDA. Of course, any agreement should be designed for the business needs of the parties. In any given scenario, some of the points in this primer will not apply, and the parties will also want to address others not included here.

This primer does not generally consider regulatory issues that may apply with respect to a JDA. Depending on the industry, the parties, and the jurisdictions involved, the parties may need to consider competition law, export control rules, data privacy regulations, product-marking requirements, laws related to government funding, regulatory approvals for products, and other compliance matters.

Contents

1. Preliminary NDAs and term sheets

1.1 NDAs

During the initial stage of a project, before a JDA is signed, the parties typically enter into a preliminary non-disclosure agreement ("NDA") to cover disclosures made during exploration and negotiation of a JDA. This initial phase is generally limited to determining each party's interest in entering into a JDA, and a preliminary NDA typically limits use of confidential information to that purpose.

An NDA should cover trade secrets and other confidential information disclosed by the parties and expressly provide that the receiving party (a) use confidential information only for permitted purposes, (b) not disclose confidential information to third parties or to persons in its organization who don't have a need to know, and (c) use appropriate levels of effort to safeguard and maintain confidentiality of information disclosed.

An NDA allows each party to comfortably disclose confidential information needed for preliminary work and negotiation. However, parties should bear in mind that the JDA may never be concluded. Accordingly, it is prudent to limit disclosure. Some information is too sensitive or valuable to share during the preliminary phase. And the party receiving information will want to be careful that it and its personnel do not become "tainted" with information from the other party that could constrain how the receiving party conducts business and develops its own technology if the JDA is not concluded.

1.2 Term sheets

Parties to a potential JDA often find it useful to prepare a preliminary term sheet or letter of intent setting out key terms of a contemplated JDA. A simplified document allows the parties to negotiate important parameters—scope of work, ownership of intellectual property (IP), license rights, cost-sharing, anticipated timing, etc.—without getting bogged down in secondary details.

Parties should be very careful to make clear that a preliminary term sheet or letter of intent is not binding and that only a final, signed agreement that has had appropriate approvals on each side will constitute a binding JDA. Even if the key terms of a deal have been set out in a term sheet, it is typical that other important terms remain to be negotiated, and it is not uncommon that parties have differences in understanding that only surface when the terms are expressed in a fully-drafted contract. If a term sheet or letter of intent becomes a binding agreement, these gaps and differences in understanding can easily lead to disputes. The required level of care goes beyond just labelling the document as "non-binding." Parties are sometimes surprised how easy it can be to create binding contractual obligations, and so the parties should also be careful in their communications and actions relating to the deal so that they do not indicate that they have a binding contract until they are ready to commit.

2. Parties to the agreement

It is a basic point that is easily overlooked—who are the correct parties to the agreement? In the context of a JDA involving one or more large enterprises, answering this question involves asking which affiliates will be conducting the work, will need rights to commercialize any resulting IP, will own IP that needs to be licensed, and have resources required to support obligations under the agreement.

For example, if a JDA participant has a subsidiary that handles manufacturing or does product distribution in the relevant jurisdiction, then that subsidiary should be included as a party, or the entity that is signing the agreement should have the right to sublicense that subsidiary to any relevant IP.

And if one party proposes that a subsidiary or other affiliate enter into the JDA, the other party should ask whether that affiliate has the resources to back-up any obligations under the JDA and the authority to grant needed licenses to IP.

3. Conduct of collaboration

3.1 Statements of work

Many JDAs contemplate multiple phases of work, and in some cases initial development will reveal the need to add scope to the JDA. As a contracting matter, it is also useful to keep business and technical details somewhat separate from the other general terms and conditions of an agreement.

For these reasons, parties should consider structuring the JDA as a master agreement with an initial "statement of work" that contains business and technical details and a process for adding additional statements of work later if the parties agree. The statement of work can include details such as:

- identification of key personnel for the project;

- technology, equipment, facilities, and other resources to be contributed by each party;

- development responsibilities and deliverables of each party;

- milestones to track progress;

- specifications and testing standards for work product to be developed;

- anticipated timelines; and

- compensation to be paid or cost-sharing details.

In negotiations, a statement of work structure makes it easier for business and technical personnel to focus on the items most relevant to their expertise and concerns, with advice from the legal team.

To avoid unexpected changes to negotiated legal terms, the JDA should provide that if there is a conflict between the terms set forth in the body of the JDA and a statement of work, the terms set forth in the body of the JDA prevail. The parties may wish to allow the statement of work to control in limited cases where the parties have made clear their intention to vary from the terms of the main body of the JDA.

3.2 Project governance

For long-term or complicated projects, parties should consider establishing some formal governance processes in the JDA. These can range from the simple to the complex.

A simple governance structure would identify principal points of contact from each party, make them responsible for overall coordination of the JDA, and in some cases give them authority to agree to adjustments in scope and timelines that will not have a material impact to the overall project or the costs borne by either party.

A more complex governance structure would add one or more committees of executives to serve strategic governance, technical review, and dispute-escalation functions. For these structures, the JDA should specify:

- roles, responsibilities, and authority of each committee;

- the seniority-level of executives from each party that will serve on the committee;

- the number of persons from each party for each committee;

- the means by which the committee makes decisions (e.g., majority vote or consensus); and

- a baseline for frequency of meetings of the committee.

Regular reporting is another governance feature that is useful in some JDAs. Particularly where the parties are doing substantial work separately from one another, they may want to include periodic progress reports as part of the JDA.

Parties should be careful not to over-engineer a governance structure in a way that adds excessive overhead to the collaboration, slows progress, or wastes executive time on matters that can be handled in the day-to-day operation of the JDA.

3.3 General standards of performance

Parties to a JDA should consider the appropriate standard of performance for their respective responsibilities.

In some cases, a JDA is intended to be exploratory or experimental, and so a general "reasonable efforts" standard is appropriate. In these JDAs, the parties intend that each will work toward the shared goals and will make specified personnel and resources available for the project, but they also acknowledge that there may be delays or that stated goals may turn out to be impractical to achieve.

In other cases, the parties will want to have defined responsibilities with more firm commitments. For example, if one party is essentially compensating the other to integrate existing technologies and both parties have confidence in the achievability of the outcome, a more definite set of commitments and time schedules would be appropriate.

3.4 Subcontracting

Some projects require that a party engage third-party subcontractors. For example, specialized fabrication or testing services may be needed for a planned product.

Where subcontracting is permitted, the JDA will generally make clear that each party is responsible for its subcontractors and also should require that any subcontractor is bound by terms that are at least as protective of the confidential information, technology, and intellectual property rights of the other party as the JDA itself. If the subcontractor will be developing intellectual property that is important to the JDA's end-results, then the party hiring the subcontractor should also be responsible for procuring sufficient rights from its subcontractor.

In some cases, the parties will want approval over subcontractors (e.g., if the subcontractor will have access to sensitive information or technology or will need to be present at a party's facilities); in other cases approval is not required.

4. Intellectual Property

The provisions regarding intellectual property rights are among the most important (and highly-negotiated) ones in a JDA. Clearly-drafted IP rights provisions help avoid later disputes and ensure that each party gets the benefit it expects from the JDA. Note that in drafting a JDA it is useful to separate and distinguish "technology"—inventions, software, works of authorship, designs, etc. and "IP rights"—patents, trade secret rights, copyrights, etc. that apply to technology. In this primer, the technical distinction is not that important, and so this primer generally refers to both as "intellectual property" or "IP."

4.1 Ownership

JDAs typically define two basic categories of intellectual property: (a) IP that each party brings to the relationship and is either in existence at signing or developed independently of the JDA ("Background IP"); and (b) IP that one or both parties create as a result of the effort under the JDA ("Developed IP" or "Foreground IP").

- **Background IP** Unless there is a specific business agreement to transfer ownership, each party to a JDA usually retains exclusive ownership of all rights in its Background IP.

  The parties may wish to identify certain Background IP in the JDA, in order to establish that particular items are owned by one party or identify IP that will be subject to a license. But for drafting purposes it is generally best to define Background IP as a category, with any listed IP as nonexclusive examples.

- **Developed IP** Under legal principles in the U.S. and many other jurisdictions, the default position is that a party will individually own the Developed IP that it solely creates. If personnel from each party collaborate in developing IP, then that IP may be jointly owned depending on the rules of inventorship or authorship for the underlying IP right (e.g., patents or copyrights). These default rules can be (and frequently are) changed by contract. Even if the parties intend to maintain ownership of IP that tracks the default legal position, it is generally preferable to specify the allocation of IP in the JDA itself. This reduces the likelihood of misunderstandings and disputes and also will help achieve a consistent approach among different jurisdictions and types of IP rights that may be relevant to the JDA.

The parties have very wide latitude to allocate ownership of IP in a JDA, and there are infinite variations that can be created. However, a few basic structures repeat frequently in actual practice and are described below.

Allocation by inventorship

Example: ToyCo and AICo.

The following sections include several illustrative figures based on the following highly-simplified example. ToyCo is a maker of children's toys, with a highly-successful line of remote-control robots with voice-recognition and other sensors and realistic-sounding speech capabilities. ToyCo is looking for its next big hit product and is interested in partnering with AICo. AICo is a developer of artificial intelligence and machine learning software for applications in various industries. The parties want to collaborate to develop a remote-control robot toy, based on ToyCo's latest design, that will incorporate a s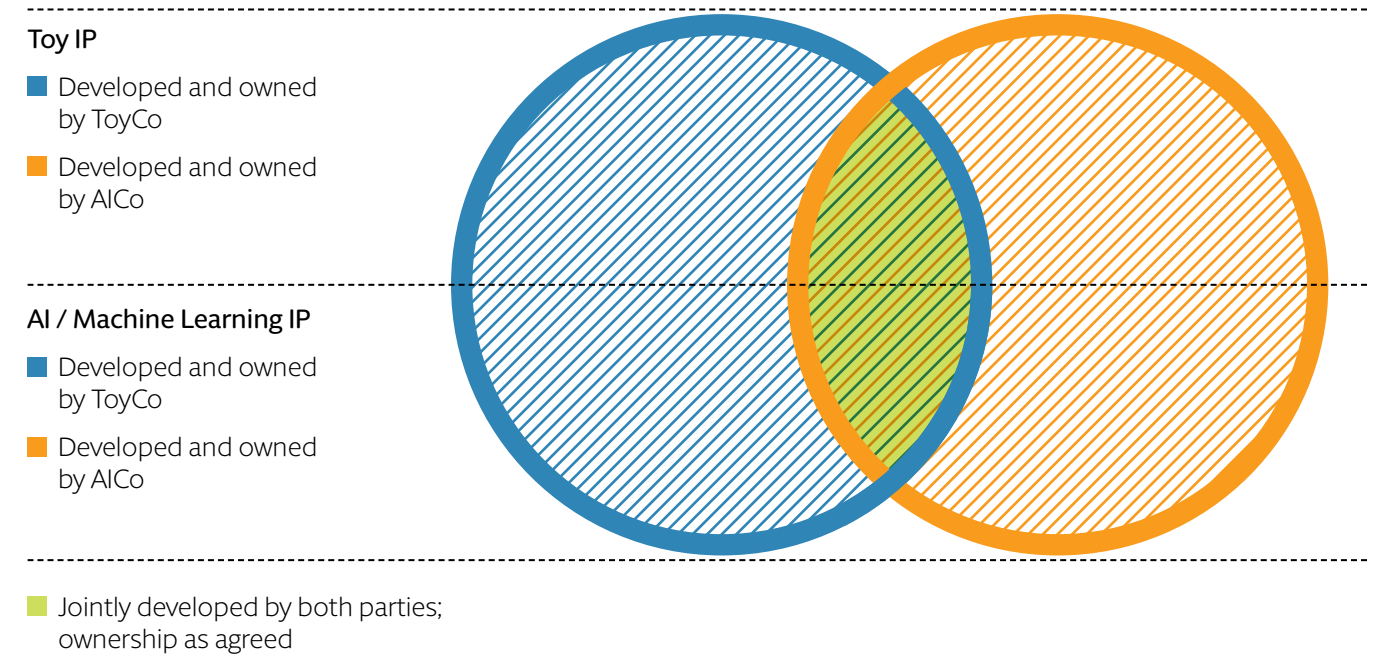implified version of AICo's software so that the robot can interact with individual members of a household, learn from those interactions, and spontaneously make the most appropriate (and amusing) comments to each person it regularly encounters. The parties are entering into a JDA for this collaboration. In general, ToyCo wants to retain ownership and control of IP related to its toy design and sensor and speech capabilities, and AICo wants to retain ownership and control of IP related to artificial intelligence and machine learning, which it plans to use for ToyCo and other customers. on its Background IP, in order to ensure that its future development is not blocked and that these patents do not fall into the hands of the first party's competitor.

The parties can agree that any Developed IP created as a result of the collaboration will be allocated based on inventorship. This means that each party owns what its personnel has developed. Allocation by inventorship is a fairly simple structure, and it also may be appealing because it seems fair to the parties—each party gets the benefit of its own work and creativity.

However, allocation by inventorship does not necessarily track the parties' fundamental business interests. It also does not account well for true collaboration between the parties, where items of Developed IP more important to one party or the other would be jointly developed and therefore jointly owned. If, for example, one party provides most of the initial technology but the other party does most of the development work, the first party will own the Background IP and the second party will own all of the Developed IP. The first party may point out that the development could not have occurred without its initial contribution. The first party may also want to control any patents relating to improvements other customers. on its Background IP, in order to ensure that its future development is not blocked and that these patents do not fall into the hands of the first party's competitor.

Figure 1 illustrates allocation of IP based on inventorship in the context of this example:

Fig. 1: Ownership by Inventorship

Toy IP

■ Developed and owned by ToyCo

■ Developed and owned by AICo

AI / Machine Learning IP

■ Developed and owned by ToyCo

■ Developed and owned by AICo

■ Jointly developed by both parties; ownership as agreed

Allocation by technology category

Another option is allocation by technology category. The parties can identify in the JDA categories that are core to each party's business. In the software field, for example, the categories could be technology that relates to an underlying operating system or middleware layer on the one hand and technology that relates to an end-user application on the other. In the semiconductor field, the categories could be technology related to semiconductor device design on one hand and technology related to manufacturing processes on the other. In this model, each party would own Developed IP that falls within the category allocated to it in the JDA.

Because it is based on technology categories that are important to each party's business, this method of allocation is more likely to correspond to the fundamental business interests of the parties. In addition, this approach consolidates ownership of related items of Developed IP, even if it was developed through collaboration.

If one party will be doing substantial development work in the other's technology category, that developing party may argue that it is unfair for the other to own the results of its effort. Of course, there are any number of ways to balance the scale of benefits for each party. A party's work may improve integration of products or facilitate use and adoption of that party's own technology in the marketplace. And negotiation of compensation in the form of engineering fees or running royalties is always an option.

When using this model, the parties must be very careful about defining the technology categories, and the potential difficulty of drawing these lines is a disadvantage of this approach. Potential overlap between categories or gaps between categories can lead to confusion and disputes.

Figure 2 illustrates allocation of ownership in the example above based on technology category:

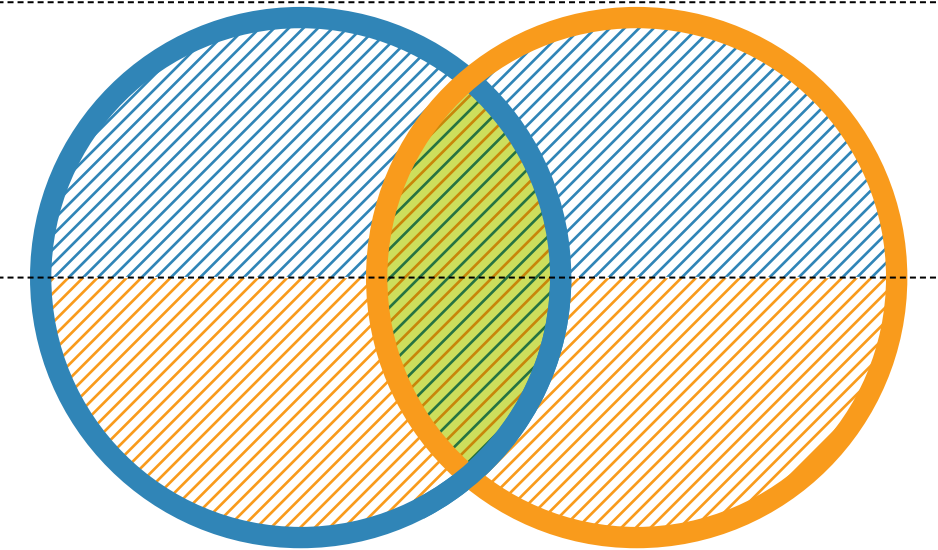## Fig. 2: Ownership by Technology Category

Toy IP
■ Owned by ToyCo

AI / Machine Learning IP
■ Owned by AICo

■ Jointly developed by both parties; ownership by category

Joint ownership

A third method of allocating Developed IP is joint ownership. In the case of a two-party Agreement, each party would have an undivided one-half interest in the whole of the Developed IP that is to be jointly owned. Parties often intend that each joint owner will have unfettered rights to use and exploit the jointly-owned IP.

On the surface, this approach seems intuitive and equitable. Each party shares equally in the benefits of their work together. But this arrangement does not necessarily track the business interests of the parties. In addition, joint ownership can get very complicated and may impact enforceability of jointly-owned IP rights—especially with respect to patent rights.

Unlike allocation of IP rights by technology category (as discussed above) a pure joint ownership approach does not allocate Developed IP to the party who is most likely to want to exploit or enforce it. Instead, both parties are tied to each other through their jointly-owned asset for as long as that asset exists.

Joint ownership can complicate use and exploitation of IP. In some jurisdictions, consent of all joint owners may be needed to license jointly-owned IP, and there may be a duty of each joint owner to account to the other and share in the proceeds of its exploitation of the IP (although these principles can typically be waived by contract).

Joint ownership can also complicate prosecution of patents on Developed IP. In deciding whether or not to seek patent protection for an invention, an IP owner has to make a choice – whether to obtain patent rights and disclose the invention to the public or maintain the invention as a trade secret. If Developed IP will be jointly owned, then the parties will need to coordinate—either in advance in the JDA or later if the decision is put off—about whether or not to seek patent protection on joint inventions.

The parties should also negotiate terms regarding which party will have control over prosecution of jointly-owned patents and how costs of prosecution and maintenance will be allocated. The parties must consider what will happen if the party controlling prosecution elects to abandon a patent or patent application. Often, that party will be required to offer to shift control (and cost) to the other party.

Finally, if a party wants to enforce rights in patents, its intentions can be impeded if the patent is jointly owned. First, if each party has the right to grant licenses to the jointly-owned IP, then each party can also undermine the other's efforts to enforce that IP by granting licenses to accused infringers. Second, under applicable standing principles, it may be necessary for all joint owners to be included as parties to an infringement lawsuit. This means that one party who has no real interest in the enforcement action, or who actively does not want to be involved, can be pulled into a dispute.

For all of these reasons, IP lawyers often try to persuade their clients to consider alternative structures that eliminate or reduce joint ownership. Happily, the parties' business objectives can often be achieved without joint ownership, through a combination of unitary ownership of Developed IP and license rights in that Developed IP.
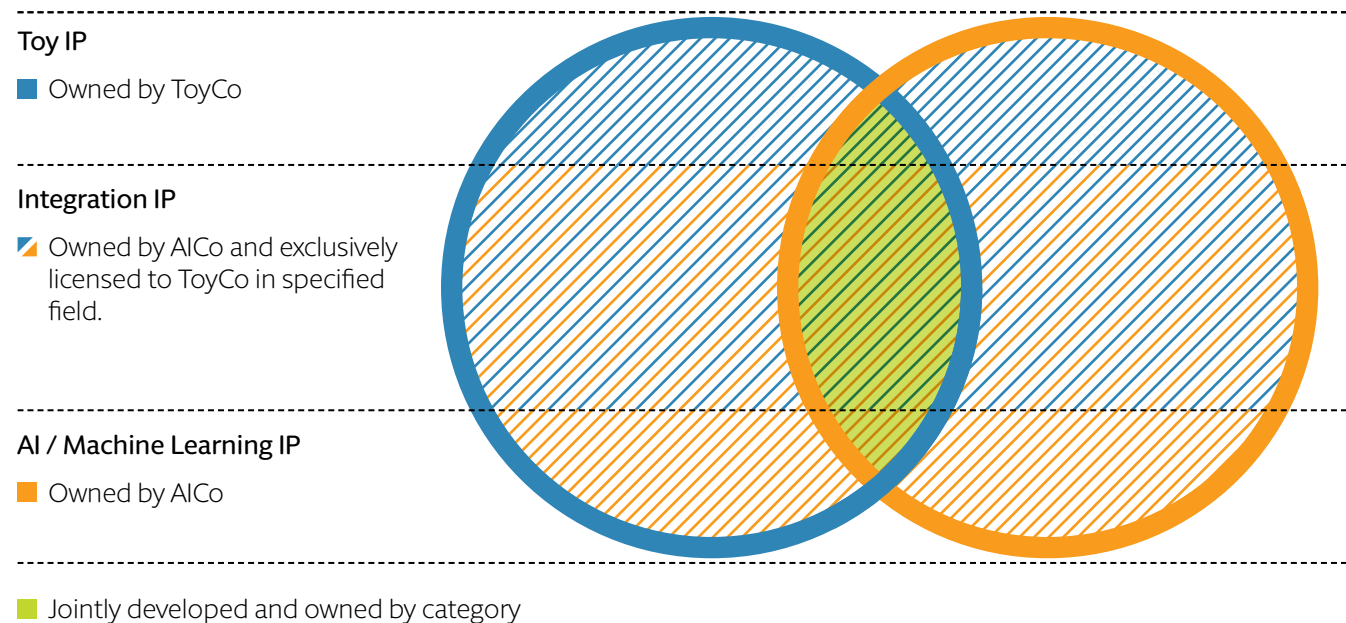
Hybrid models

In real practice, many JDAs adopt hybrids of the three models described above, with some Developed IP allocated by category and Developed IP outside of the core categories allocated by inventorship, for example.

Two keys for success are maintaining clarity about which Developed IP goes in which categories and keeping the agreement simple enough that it can be effectively understood and administered.

Figure 3 illustrates a possible hybrid ownership model for the example above, with the addition of a "middle" category of technology that is related to integration of device IP and manufacturing IP:

Fig. 3: Potential Hybrid Model

Toy IP
■ Owned by ToyCo

Integration IP
◢ Owned by AICo and exclusively licensed to ToyCo in specified field.

AI / Machine Learning IP
■ Owned by AICo

■ Jointly developed and owned by category

4.2 Licensing

In addition to allocating ownership of IP, parties to a JDA must consider what licenses to both Background IP and Developed IP should be granted. There are a number of kinds of licenses that recur in JDAs.

First, if the parties will be collaborating in development efforts, each party will typically grant the other a license to use the Background IP that is furnished under the JDA. This license will be limited to use for the collaborative efforts under the JDA and will end when the JDA (or an SOW) ends. Depending on the circumstances, the parties may wish to have a blanket license to Background IP that is made available or may wish to be very specific about which Background IP is licensed or not licensed. This Background IP license usually does not include commercialization rights. The limited duration and scope of this kind of license makes it easier for a party to grant, but each party should also consider additional limitations and restrictions, including confidentiality restrictions, that should apply to Background IP.

If a JDA provides for Developed IP to be transferred from one party to the other, or if one party will have sole ownership of Developed IP that is jointly developed, then the parties may agree that the party that does not own this Developed IP will have a license. The scope of the license (e.g., for research and development only or also for commercialization) will be negotiated depending on the business circumstances. For example, in the toy robot example, one party has a strong interest in ownership of all sensor and speech-generation IP that is developed, but the parties may also agree that the other party will have a license to use some of that IP that is not specific to the robot toy prodcut for use in AI and machine-learning applications generally.

In some JDAs, the parties will also agree on licenses to cover commercialization of Developed IP that results from the project and possibly Background IP on which it depends. In other cases, the parties may want to focus the JDA on development and defer discussions of any commercialization license. It is often difficult to know in advance which party will have leverage in license negotiations once the development work is complete, and so each party should consider carefully whether or not to defer this discussion.

For each commercialization license, key terms to be negotiated include:

- field of use, products, and services licensed;
- scope (use, make, have made, sell, etc.);
- sublicensing rights (if any);
- royalties (if any) and related terms;
- exclusivity (if any), including scope, duration and conditions (such as minimum royalties);
- duration;
- termination rights; and
- transferability and effects on the license if there is a change of control of the licensee party.

4.3 Right of first refusal

If Developed IP is allocated to one party or is jointly owned, the parties may want to provide that the other has a right of first refusal, right to negotiate, or other preferential position if the other party desires to transfer or grant exclusive rights under its interest in the Developed IP. The provision governing these rights should clearly define what constitutes a transfer opportunity that will give rise to the right. It should also delineate a process for party to exercise its rights or not. Any right of first refusal or similar right will make it more difficult for the owning party to dispose of the Developed IP it owns, and so the owning party will want to make the process, and its endpoint, very clear.

5. Confidentiality

In a JDA, the parties will typically be sharing confidential information and will seek to protect that confidential information with a confidentiality provision or separate NDA. While sometimes overlooked as a "boilerplate" provision, the confidentiality terms of a JDA deserve careful attention so that each party's trade secrets are protected and the intended allocation of IP ownership and license rights is not thwarted.

Where the parties have agreed to a Preliminary NDA, they may wish to incorporate that existing Preliminary NDA into the JDA. This has the advantages of streamlining the JDA somewhat and also maintains continuity of protection of confidential information. However, the parties should review the Preliminary NDA again to make sure that the permitted uses, permitted disclosures, definitions of confidential information, duration of protection, and other key terms will work for the longer-term JDA and amend or replace the Preliminary NDA if they do not.

Where the parties want to include a confidentiality provision in the JDA, they should consider the following:

- definition of confidential information (e.g., must information be marked confidential, are certain categories always deemed confidential, etc.);

- permitted uses, including uses as permitted under licenses granted in the JDA;

- duration of protection;

- circumstances in which protection ends (e.g., if information becomes public without fault of the receiving party);

- exceptions for legally-required disclosures; and

- exceptions for disclosure in patent applications that are filed in accordance with the agreement.

The parties may also want to include a separate provision governing disclosure of the existence and terms of the JDA itself.

If one of the parties to a JDA is an academic institution, the institution likely has a strong interest in maintaining the ability to publish the results of the collaboration. In this case, parties will often agree on a process that allows the other party to review proposed publications, make appropriate requests for continued secrecy, and file patent applications prior to publication where appropriate.

6. Term of the agreement

The parties should use care in defining the duration of the JDA. Normally, each party will want to put a time limit on its responsibilities under the JDA. But neither party will want it to end prematurely—before the parties have had adequate opportunity to pursue their goals or before the expected duration of any license rights being granted.

Particularly in a JDA with multiple statements of work, it is often useful to define a term for the collaboration under each statement of work and a separate, longer term for any commercialization licenses that are granted for the results of the JDA. That way, the active development part of the JDA can end and the "next phase" under the licenses can continue.

A JDA also typically defines termination rights for each party. Termination rights for uncured material breach are typical. Termination based on failure to meet defined development milestones is also common. The parties may want to permit termination for convenience as well—so that a party can exit without having to show cause. If termination for convenience rights are included, the parties will need to consider what limitations are appropriate, including for example notice periods and a minimum duration before notice can be given. Parties should also consider what happens upon each kind of termination—whether compensation or cost-sharing rules should apply, what licenses and other provisions should survive termination, etc.

7. Other terms

This primer has focused on several key terms that are somewhat unique to JDAs. Of course, parties to a JDA will need to consider a number of other terms and conditions, including the following:

- allocation of costs for development activity and facilities, equipment, and materials used;

- taxes;

- export control and compliance with laws;

- representations and warranties;

- indemnification for third party claims (including claims relating to intellectual property infringement, product liability, and breaches of key terms of the JDA);

- limitations of liability;

- data security and privacy;

- assignment and delegation;

- term of the agreement; and

- termination and effects of termination on licenses granted.

Alicante

Amsterdam

Baltimore

Beijing

Birmingham

Boston

Brussels

Budapest

Colorado Springs

Denver

Dubai

Dusseldorf

Frankfurt

Hamburg

Hanoi

Ho Chi Minh City

Hong Kong

Houston

Jakarta

Johannesburg

London

Los Angeles

Louisville

Luxembourg

Madrid

Mexico City

Miami

Milan

Minneapolis

Monterrey

Moscow

Munich

New York

Northern Virginia

Paris

Perth

Philadelphia

Rio de Janeiro

Rome

San Francisco

São Paulo

Shanghai

Shanghai FTZ

Silicon Valley

Singapore

Sydney

Tokyo

Ulaanbaatar

Warsaw

Washington, D.C.

Zagreb

Our offices

Associated offices

www.hoganlovells.com