

The business value of open collaboration

How interoperability brings manufacturers and consumers together to build a smarter planet



Introduction

Building a smarter planet demands software everywhere. Software runs our business processes, back-office functions and customer transactions. It drives our web presence, online business and customer relationships. And increasingly, it's embedded in the smart physical products and infrastructure of everyday life—from automobiles, consumer appliances, phones and personal devices, to power grids, traffic control and healthcare systems. Everything we use is increasingly made smarter by software and connected through the Internet.

Consequently, the ability to rapidly and reliably enhance enterprise software systems and the ability to deliver smarter products and services have become critical differentiators for businesses that face increasing competitive pressures. To meet these challenges, business leaders need their software delivery process to work at least as well as their other business processes. They need to have predictable delivery cycles; cost-to-value management; and continuous visibility into work in progress. They need to be able to rely on the delivery of the software elements of business initiatives and products—but all too often, despite the best efforts of vendors and users, they can't.

Even though development and delivery team members are smart, hard-working and dedicated to satisfying their stakeholders, somehow problems still occur. Requirements get lost in translation from analysis to construction. Test cases get missed, causing old bugs to reappear. Project progress is hard to decipher and status is unreliable. Even in the most successful, well-run organizations, too much effort is wasted on mundane, administrative tasks such as data transfers or status updates. Conversely, overly-prescriptive processes risk stifling both creativity and productivity.

Why is it so hard for smart, hard-working, dedicated professionals to successfully deliver software projects? One major barrier is often the tools that software development professionals depend

on. More accurately, it lies *between* those tools. The software delivery process today is reminiscent of a 1950s office. The simple act of sending a letter might involve the following steps:

- A manager dictates the letter to a secretary, who takes it down in shorthand.
- The secretary delivers the shorthand to the typing pool.
- A typist transcribes it, making a carbon copy, and sends it back for review.
- The manager reads it over and marks any corrections.
- The secretary takes it back, and the wasted carbon copy is discarded.
- This cycle repeats until the letter is satisfactory.
- The manager signs it and the secretary mails it.
- Finally, the secretary files the carbon copy in case it is ever referred to again.

Just imagine the number of ways that this process can go wrong. The secretary mishears the dictation. The typist misreads the scribbled shorthand. Or the manager realizes he omitted critical information when he reads back his own letter—to name just a few. Each handoff introduces the possibility for errors and omissions, causing costly rework and delays, not to mention the expense of verification required at every stage to make the process go smoothly. Unfortunately, this quaint last-century vignette is all too reminiscent of the process handoffs in modern software development. In the case of software development, the problem is even greater as development is far more complex than the simple task of correctly transcribing a letter. It is an inherently creative and exploratory endeavor, filled with uncertainty, ambiguity and unpredictability.

In this paper, you will learn the value of “openness” in business computing, including open standards, open source and open interfaces. We will also describe an initiative that breaks down the

barriers between tools so that teams can collaborate seamlessly, communicate effectively and provide the visibility that the business needs—without requiring companies to retool and reskill their development organizations.

It takes more than a village

Throughout the history of the software development industry, vendors have created increasingly helpful and sophisticated tools that have dramatically increased the productivity, accuracy and effectiveness of every role in the software development life cycle. But often the effect has been as if we replaced those 1950s typewriters with word processors, and left every other element of the process in place. Individuals have better tools today than they have ever had, but the process itself is still flawed. No matter how good the business analyst's requirements tool is, effort will be wasted if she has to manually hand off those requirements to designers (as if she were an executive dictating to a shorthand-recording secretary). It doesn't matter how good the test automation software is if the defects get passed back to development unreliably (as with hand-scribbled notes on a first draft of a letter). Development organizations need a solution to the challenges that face the *team* across the life cycle, rather than merely the challenges facing each *individual*. What they need is collaboration.

In an attempt to meet this need, the industry has seen the rise of successful vendors offering cross-life-cycle tools, such as those provided under the IBM® Rational® software brand, that provide the necessary integration and interoperability. However, it's rare that a software delivery organization relies solely on a few tools from a single vendor. More often, teams bring to bear a range of commercial, open source and in-house developed tools. Often they employ competing tools in different parts of their organizations. Realistically then, this discussion can't be about a specific tool suite or small number of vendors. Moving beyond the 1950s analogy requires an industry-wide perspective.

The industry has made numerous attempts to solve the burning need for tools interoperability, but each of them foundered for some combination of the following reasons:

- The scope was too ambitious. Any attempt to define a comprehensive “integration map” for the whole software life cycle was doomed to be obsolete before it could be finished.
- It depended on unrealistic levels of vendor cooperation. Each vendor involved in standardization efforts would try to entrench its own proprietary advantages in the standard, and each would just as fiercely resist attempts by its competitors to do the same.
- It required “rip-and-replace” adoption by user companies. Under some proposals, customers would have to replace their existing tools with new ones. At the very least, they would need to migrate data to a single common repository. This was too high a barrier for potential adopters.
- The collaboration was not sufficiently open. Often, the invitation to participate in a standardization effort was driven by competitive motivations rather than even-handed inclusiveness, leading to the omission of key potential contributors and a lack of credibility.
- The focus was on what vendors wanted, not on what their customers needed. As a result, protectionist compromises limited the business value that user companies could obtain.

Lacking any life cycle-wide strategy for integration, vendors have resorted to the next-best alternative: individual bridges between pairs of tools, using vendors' published proprietary APIs. The disadvantages of this approach are many, and include:

- The number of integrations rises exponentially with the number of tools. Developing and maintaining the integrations is expensive for vendors, and operating them is burdensome for customers.
- Coverage is limited. Each vendor tends to integrate with the tools of its allies, and ignore or exclude the tools of its competitors, to the detriment of customers who would simply like to choose the best tool for each job.

- The integrations are “brittle.” Pairs of tools become tightly dependent on each other’s internal structures or detailed behaviors, and upgrading one product can break the integration. And worse, since each of those tools is on its own release schedule, customers using a variety of tools may struggle to find a mix of releases of the interdependent products that work together.

The greatest integration on earth

Faced with this pressing business need and historical difficulty, the IBM Rational organization took a step back and asked: What kind of integration approach would it take to make life cycle collaboration successful? We decided that the architecture of the solution would have to be:

- **Incrementally adoptable by user companies.** It must not require wholesale replacement of existing tools or skills or mass migration of data. Instead, it must allow a step-by-step implementation, with each step delivering positive value.
- **Technically acceptable to vendor companies.** It must not mandate a specific implementation technology or repository that favors one vendor over another. It must not require vendors to make wholesale changes to their existing products or give up proprietary product differentiators.
- **Authentically open.** Participation would have to be open to any vendor, open source project, or in-house implementer, regardless of competitive positions or alliances.

As we looked around for a successful means of integration that fit this profile, we realized that the best precedent is also the largest example of “integration” that the world has ever seen: the world wide web itself. By conforming to a modest set of standard interfaces and a common architectural principle, anyone can participate in the web by creating a web site or publishing a blog; and any browser can consume any of the billions of pages that

make up the web.¹ Many different types of content coexist, and new types are easily added without disruption. Furthermore, the simple and flexible architecture of the web, and in particular its common integration mechanism of resource links, allows the emergence of new uses for all that content that did not exist when the web was first invented, such as search engines or content aggregation.

On the web, no one needs permission to create content, or to consume content, or to link to content.² The result of this open access is that Internet participation has grown exponentially, while membership in older-style online services that restricted content and limited users’ choices has rapidly declined.

At IBM, we realized that the way our industry has approached integration in the past is like those closed, limited online services. We asked ourselves: how could software life cycle collaboration be more like the open, unrestricted web? The answer we propose is agreement on basic specifications for sharing life-cycle resources. The vehicle to get there is a forum for cooperation called Open Services for Lifecycle Collaboration (OSLC).

OSLC: An open forum for collaboration

The goal of OSLC is to re-examine how we look at interoperability so that:

- Any tool can be integrated on an equal footing simply by sharing resources and services agreed to in open specifications.
- Anyone can participate in the specification process on equal terms, not only vendors but also user companies, industry forums, open source projects, or motivated individuals.
- Any organization can freely take advantage of the specifications, whether they choose to contribute to the effort or not.
- The mistakes of the past that led to limited coverage and brittle integrations are averted.

To achieve those ambitions, OSLC embodies three key elements:

- A unifying and universal architectural style that takes advantage of the lessons of the web. That style is what the web community calls *linked data*. Linked data was conceived as a set of best practices for using the web to connect related data that wasn't previously linked, or using the web to lower the barriers to linking data currently linked using other methods. OSLC is inspired by the idea that what works for data on the web can also work for life-cycle data.³
- A set of technical specifications inspired by real-world scenarios, openly published and freely adoptable. This means creating specifications that apply linked data concepts to the domain of life-cycle tooling, or simply, specifications for linked *life-cycle* data.
- A transparent and open community process with no unreasonable barriers to participation.

The OSLC organization today is simply a group of people who share an interest in changing the status quo. They use modern on-line collaboration tools for discussing specific integration challenges and scenarios. Solutions are proposed, debated and prototyped, and ultimately, candidate specifications are published. OSLC has no membership fees or applications; no egregious bureaucratic processes; and no “purity tests” for participation. In fact, there is only one way to influence the work of the OSLC—to participate in it.

As we established OSLC, we knew that it would be critical that the intellectual property required to implement compliant products be open. However, it turns out that there is more than one way to be “open,” and to choose the appropriate way, we had to take a step back and review the business goals we were hoping to achieve.

The business value of open markets

Since the beginning of the industrial age, smart companies have recognized the value of common specifications. Whether its house bricks, screw threads, car parts or electronic circuitry, markets work best when suppliers, purchasers and users can rely on the interoperability and interchangeability of parts regardless of source. Shared specifications promote competition, diversity of offerings, reuse of common components and innovation, because incremental improvements can be built on top of the existing foundations. Everyone is better off as a result: Sellers create more value, buyers receive more value and innovators can add more value. The compact disc (CD) market flourished because Sony and Philips, the vendors who created the technology, licensed the necessary patents widely, reasoning correctly that both were better off sharing a huge market than owning a small one. By contrast, the adoption of both HD-DVD and Blu-Ray Disc, competing formats to succeed the DVD, was stalled for several years by neither specification becoming universally endorsed.⁴

And yet, for much of its early history, the IT industry shunned cooperation, with vendors preferring to protect proprietary fiefdoms rather than promote collaboration. For many years, the word “standard” in IT was doublespeak for “dominant proprietary product.”⁵ This practice continued despite the growing evidence that common specifications were better for everyone. Buyers benefit because vendors are forced to compete on quality, price, support and pace of innovation. Vendors benefit because the market as a whole grows more rapidly—so even though they have to compete with each other, the amount of business value available for them to target more than compensates for the loss of a small, protected niche.⁶

Less obvious but equally important: open markets create the opportunity for an *ecosystem* to emerge in which innovators can create compatible added-value add-ons, each finding their own

niche. For example, the definition of standardized slots and connectors in car dashboards provides customers with an enormous choice of entertainment options, both directly from the car manufacturer or as an “after market” upgrade. And of course, a car owner can upgrade the radio to replace the tape player with a CD player or the CD player with an iPod dock at any time, without having to replace the whole car! Even as the entertainment technology has advanced, the interface specification has remained stable.⁷ In fact, the creation of a value-enhancing ecosystem that allows many participants to benefit is often the single most important predictor for widespread adoption of a standard for technology.

There are numerous proof points in IT for both the benefits of open specifications and the importance of an ecosystem. The database market boomed with the adoption of the SQL open standard, in part because a variety of vendors could offer value-added tools for monitoring, tuning and so on, and because skills such as database design became more portable. The personal computer market grew dramatically when IBM established the PC as a dominant format, which allowed:

- Software authors to target one common base of customers rather than a dozen fragmented markets.
- Peripheral manufacturers to consolidate their costs around a single technical standard.

More recently, the establishment of the Open Document Format (ODF) has allowed many different vendors to offer word processors and office suites in competition with Microsoft, whose proprietary format used to determine the “standard” in that market. Many user organizations are mandating ODF because an open standard guarantees them the crucial ability to access their own documents in the future without being governed by the decisions of one company.

Today, the IT industry has, for the most part, embraced open markets, and it is rare to see a new technology initiative launch without a corresponding roadmap for establishing an open specification. However, as the examples above show, “open” comes in various flavors, ranging from a dominant proprietary format that is widely licensed to a formal standard approved by an independent international body. It’s important to choose the right one for the business objectives at hand. The question that must be asked—and that is all too often is overlooked—is simply: What must be standardized to increase business value and decrease business cost?

Open standards

The most common meaning of “open” concerns itself with standardization of the *behavior* of software. For example, the purchaser of a standards-compliant SQL database or a JEE application server can be confident that the software will produce the same essential results, regardless of the supplier. Of course, each vendor is still free to compete on quality, reliability, cost, support, and extensions of functionality beyond the standard; and to patent or keep as trade secrets uniquely effective methods of implementing that standard behavior. Thus, the benefits of open standards can be compared to the many standards for implementation that are commonplace in industry and engineering, such as DIN, ISO, ANSI and more. Open standards gained broad adoption in IT during the 1980s, a period that saw the emergence of cross-industry organizations such as the OMG, Open/X and the IETF.⁸

Open source

The 1990s saw the commercial rise of a new meaning of “open”: open source.⁹ In this model, the source code for an application is made freely available to those who wish to use it, subject to certain restrictions (depending on the specific license used to publish the source code). The most common restriction is the requirement that any modifications to the source code must be

shared back with other users under the same terms. Unlike open standards, which promote common *behavior*, open source promotes common *implementations*. There are few parallels outside of software to the value of open source, because of the unique nature of software. One loose analogy is to classic books that are out of copyright. The text of such books might be available in numerous editions from different publishers, including hardback, paperback, annotated, collected with other works and so on. But even this analogy fails to capture the true value of open source software, much of which comes from the unique flexibility with which software can be modified, adapted and extended. In fact, some of the most commercially significant open source projects have been those that promote a common “platform” allowing innovators to add value without having to constantly recreate the common parts. Furthermore, it allows each innovation to enhance the value of the others. Individual innovations may themselves be open source, or use conventional commercial licenses.¹⁰

Open interfaces

Most recently, a new meaning for “open” has gained prominence: open interfaces. This model came about from the realization that often, companies want to integrate a wide variety of tools or processes that need to be able to exchange information, even though each may not know or care what the other endpoint does with that information. To understand how powerful this idea is, think of it as analogous to the electricity supply. You don’t care how the electricity was generated, whether by wind, coal or nuclear power. You don’t care what path it took to get to your office or what voltage transformations it went through on the way. You only care that when it reaches your power receptacle, it provides the expected voltage and frequency. Conversely, the power company doesn’t care what kind of appliance you plug into your sockets. Open interfaces in IT provide the means for applications to act like “utility suppliers” of information to each other.

One of the earliest and most widely known successes for this model was in the arena of news “syndication.” A simple standard¹¹ was established that allowed servers to publish a series of changes as a “feed,” often as news headlines and associated stories, and allowed clients to aggregate many different feeds and present them to a user. The servers are free to choose any implementation, frequency of publishing, and format of their contents that they want, as long as the *interface* follows the standard. Conversely, users can choose from a huge variety of ways to view, organize and filter their chosen feeds. Today, a vast array of information sources—far beyond traditional news stories and ranging from personal blogs to system maintenance logs—can be consumed as feeds.

The openness of OSLC

Faced with all these choices, we had to ask ourselves: what kind of “open” should OSLC be? As we considered the business drivers, we concluded:

- Demanding that companies reskill and retool, or that participating vendors substantially redevelop their existing products, was not an option. Therefore, the common implementation approach implied by open source was not appropriate.
- Software development and delivery involves not only many different tools but also many different approaches. So standardizing the behavior of even the most commonly used types of tools would make little difference to the overall challenge. Therefore, the common behavior approach implied by traditional open standards was not appropriate.
- To support the diverse array of tools, processes and business needs—including many we don’t even yet know about—we realized that the most important thing to agree on was the interfaces between different tools.

By focusing on agreement about the interfaces, we could allow each tool provider to participate without having to predetermine all the ways that users might integrate their particular tool, or requiring close relationships with all the other tool vendors a user might choose. We could also allow each user to integrate their choice of tools in the way that made the most sense and delivered the most business value in their own circumstances.

We also determined that OSLC interfaces had to satisfy a number of rigorous requirements, including the following:

- **Neutral.** The interfaces must be technology-neutral, and could not favor the implementation details of any particular vendor's existing product.
- **Adaptable.** It must be feasible to retrofit OSLC interfaces to existing products, without extensive reengineering of products currently in use.
- **Universal.** This was the most critical and most challenging requirement. It means that any asset, resource or artifact used in the software development and delivery life cycle must be able to reference and associate itself with any other, regardless of type, location or implementation.

As noted earlier, as we considered these requirements, we realized that there is already an outstanding example that satisfies all of these needs—the web. All links are URLs (universal addresses), and they have the critical benefit that any web page or online resource can point to any other, without depending on any knowledge of what is at the other end of the link. In OSLC, we have adopted exactly the same linked data approach as the web itself. OSLC specifications depend on a small set of universal principles:

- Any asset can refer to any other asset using just one mechanism, namely a URL that identifies the location of that asset. Conversely, to make its assets available to any other OSLC-compliant tool, a tool need only expose an appropriate URL for each asset.

- The list of operations that one tool can perform against another is very short and is also universal. For example, tools can request content from each other, send updates to content and follow links from one resource to another in a standardized way. If another tool that provides the same type of resources and the same set of simple operations is substituted, the integration should not change.
- The exposed format, or representation, of a resource is independent of the internal format, technology or representation of any particular provider. Furthermore, the requesting tool can ignore any part of the representation that it doesn't understand or isn't concerned with (provided it preserves the information when it passes it on or passes it back).

When tools expose OSLC interfaces, the assets that they house can be linked and used just like a “web” of software development and delivery resources. One powerful consequence of this is that OSLC integrations are “loosely coupled” and independent of the target tool. For example, if one part of your testing organization is using Tool A from one vendor while another part is using Tool B from another vendor, your requirements analysts can link their requirements to test cases in exactly the same way, regardless of the test team with which they collaborate. And if the requirements analysts decide to switch to a different requirements tool, it does not disrupt their collaboration with the testers, provided that the new tool also implements the OSLC interfaces.

The OSLC community

The community approach that we are taking at OSLC is just as critical to its success as the technology neutrality described above. We recognize that customers frequently mix tools from several vendors, as well as open source projects and “home-brewed” utilities. Experience has shown that any vendor-specified set of APIs invariably privilege one vendor's tools and disadvantage others'—or at the very least, make other vendors suspect this is the case. Vendor-specified APIs also leave many business partner companies struggling to catch up as the vendor evolves the API independently or in collaboration only with its

closest partners, leading to suspicion and reluctance to participate by competing vendors. The result has been the fractured and unsatisfactory integration landscape of the past, as described at the beginning of this paper.

Therefore, to promote community involvement and ensure transparency, all of the work of OSLC is taking place in the open, at <http://open-services.net>. There, OSLC is organized by domain-specific working groups in areas such as change management, requirements management, quality management and software configuration management. Each work group defines and prioritizes the scenarios, describes the scope of each iteration and writes the specifications for interacting with tools in their domain of interest.

So that the overall objectives of OSLC are met within each working group—and to avoid the creation of silos that do not interoperate—the project leads from the OSLC working groups coordinate their work across the domain areas. The project leads also collaborate on topics of common interest and in sharing design best practices that they discover in the course of their work.

To help ensure that the content developed for OSLC is broadly consumable, OSLC has adopted two important intellectual property policies. First, all content posted to the site is covered under a Creative Commons license (See Appendix B.), with liberal usage rights. This license allows anyone to freely use the specifications without any license fee or restrictive agreement. Companies don't even have to be members of OSLC to use the specifications. They are freely available to everyone.

Second, all contributors to an OSLC specification publish a patent non-assert covenant, promising not to assert any “necessary claims” against implementations of the specifications, or grant a specific license to implementers of a specification. These

policies have been important in securing the participation of both open source and commercial concerns. OSLC participants include individuals from software vendors, open source projects, systems integrators, industry IT teams and the academic community.¹²

Conclusion

Overcoming the challenges of tool interoperability isn't easy, but it is important to the teams that apply tools as they try to differentiate their businesses. It's important to tool providers who face demands for an ever-increasing set of pair-wise integrations. And we would argue that it's important to the industry as a means of breaking software delivery resources out of the IT silo and making them more accessible and interoperable with the broader business processes. OSLC seeks to change the game, but in a practical way, learning from the lessons of the past and building on the success of the Internet.

One thing is clear. OSLC gets better with participation. The more people get involved, the better chance we have to establish real collaboration across the life cycle. If you have something to contribute, we encourage you to join at <http://open-services.net>.

Appendix A: The technical architecture of OSLC

To collaborate across the software life cycle, the most basic technical capability is the need to create and manage navigable relationships between “resources” or assets. Resources include everything from requirements to code to test cases to configurations. Complicating matters, some of these resources can be complex formatted documents such as a requirements specification; others are collections of other resources, such as a configuration; others are artifacts created by the process itself, such as a build log. All of these different resources need to be treated in a uniform way if tools are to be able to link their assets to each other without hard-coded foreknowledge of the target endpoint.

Identifiable and addressable: The value of URL-accessible data

On the web, this problem has been solved. All links are URLs (universal addresses), and they have the crucial benefit that any web page or resource can point to any other, without depending on any knowledge of what is at the other end of the link. In OSLC, we wanted to create a similar means of relating resources that, like the web, satisfied three basic criteria:

- It must be vendor- and technology-independent. It must provide a way for any vendor's representation of a test case, for instance, to point to any other vendor's requirement.
- It must be location-independent. Resources are scattered across multiple repositories and tools, possibly widely or even globally dispersed. And those resources can move from time to time.
- It must be able to solve the problems we don't yet know about, not just the ones we do. The future will certainly bring new types of assets into the life cycle, as well as new styles of integration and aggregation. Like the web itself, the linking mechanism must be flexible enough to adapt to change.

Identification and location are the most fundamental parts of any digitally based interoperability story. Like the web, every resource in OSLC has a unique address that tells any other tool how to locate it—a URL. It is impossible to overstate the fundamental importance of having URL-accessible data. This is what makes an application's data *identifiable and addressable* from any other application, in a completely technologically neutral fashion. A tool that needs to refer to a resource used by another tool, such as a test tool with a test case that references a requirement, need store only the URL that identifies that resource, just like one web page linking to another. By contrast, many existing Application Lifecycle Management (ALM) tools rely on proprietary naming schemes that require deep knowledge of the tool's technology to resolve.¹³

Having a uniform way to identify and locate resources is essential, but by itself it merely allows us to establish pointers or references from one resource to another. To create more useful integrations, we need one more thing: a standardized “interface,” or common set of services, that are available for resources. Again, the web has provided an answer for us that we use in OSLC: a small set of actions that can be used identically from any tool to any other tool, regardless of the type of resource or the technical implementation. OSLC uses a style of integration called Representational State Transfer (REST) to fetch or modify any resource, regardless of its type or location, ensuring that tools remain independent of each other's implementation details. In a RESTful architecture, clients initiate requests to servers using a small set of general purpose services. Servers process those requests and return appropriate responses.¹⁴

Representations

In OSLC as on the web, the notion of a resource is very general. It could be a requirement, a work item, or even a web page that enables a user to create or select from a list of resources that meet certain criteria. These are all resources, identified by their URLs.

URLs, along with a RESTful architecture, allow us to locate and access such a resource, which is often useful by itself, but, by design, these URLs do not provide a tool with any information about what's “inside” the resource. The content is still unknown. For example, a requirement could be represented by a text document describing the requirement, an image showing a screen mockup, an XML document defining the attributes of the requirement, or any one of many other representations. A test case management tool that wants to verify the existence of a test case for every requirement need not understand the requirement's contents. It need only confirm its existence and location (i.e., its URL). When the tool's users want to see or edit the contents of the requirement, it functions exactly like a browser. It navigates the link to retrieve the resource and hands it to an appropriate tool that understands the resource type.¹⁵

Although we value this flexibility, we can do more with the resources when we know some details about their format. Therefore, much of the work in OSLC working groups is to define agreed-upon resource representations. With this information, any tool can examine the common elements of these resources. This allows much deeper integration between tools.

For example, a tool could view and modify the description of any life-cycle resource. A quality management tool might refer to a requirement stored in a requirements management tool and be able to reflect whether associated test cases pass or fail. A configuration management tool might flag tests stored in a test case management tool as needing to be run during the next regression test because relevant code modules have been changed. A management console might pull information from a wide variety of tools to produce a consolidated view of project progress and status.

And, just like the web, the retrieved representation of a resource might contain links to further resources. For example, a defect implicates a code module that is tested by a test case that satisfies a requirement. Each of these links exploits the same mechanisms: URLs and uniform services.

Appendix B: Creative Commons licenses and copyrights

OSLC chose to publish its specifications under Creative Commons licenses so that everyone could adopt these specifications with confidence. A Creative Commons license is a mechanism that allows a copyright holder, such as the authors of an OSLC specification, to renounce certain rights that they would normally have under copyright law, while retaining others.

This mechanism allows any implementer to take and use the published OSLC specifications without risk that some kind of license claim will be enforced against them.

Critically, a Creative Commons license cannot be revoked. For example, if a vendor implements the published Change Management 1.0 specification, the copyright owners can *never* revoke that vendor's permission to use that specification in the ways initially permitted.

More information about Creative Commons is online at <http://creativecommons.org>

For more information

To learn more about the business value of open collaboration, please contact your IBM marketing representative or IBM Business Partner, or visit the following website: <http://open-services.net>

Additionally, financing solutions from IBM Global Financing can enable effective cash management, protection from technology obsolescence, improved total cost of ownership and return on investment. Also, our Global Asset Recovery Services help address environmental concerns with new, more energy-efficient solutions. For more information on IBM Global Financing, visit: ibm.com/financing

About the authors

Carl Zetie: Strategy team for Rational brand; Market Segment Manager, IBM

Scott Bosworth: Rational CTO Team; Business Programs Manager – Operations, IBM



© Copyright IBM Corporation 2011

IBM Corporation Software Group
Route 100
Somers, NY 10589 USA

Produced in the United States of America
May 2011
All Rights Reserved

IBM, the IBM logo, ibm.com and Rational are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

Other product, company or service names may be trademarks or service marks of others.

⁷ In the US at least, the cassette deck is now a thing of the past. The last car available with a factory-installed cassette deck in the US was reportedly a 2010 Lexus. Because of the standardization of interfaces, individuals can, of course, still add a cassette deck from the wide choice of aftermarket units, just as they can replace the cassette decks in older cars with CD players.

⁸ Some basic standards have existed from the earliest days of IT, especially in areas strongly influenced by IT’s highly standardized counterpart, telecommunications. For instance, ASCII dates back to the 1960s, and other early examples include the programming languages COBOL and FORTRAN. However, it was not until the 1980s that higher levels of system behavior were widely standardized.

⁹ The open source movement had been influential in academic circles for at least two decades before it began to impact commercial software adoption, for example at MIT in the GNU project. However, the impact on commercial IT rose dramatically with the creation of Linux in the early 1990s and the Apache web server in the mid 1990s.

¹⁰ One highly successful example of this model is the Eclipse Foundation. See <http://eclipse.org/>

¹¹ There are two competing standards that are widely used today: Really Simple Syndication (RSS) and Atom. Although they differ in detail, both serve the purpose described here and both share the distinction of standardizing only the interface between servers and clients. The name RSS is often used casually to include both standards.

¹² Learn more about OSLC community members at <http://open-services.net/html/Snapshot.html>

¹³ For example, one tool may require you to know the case-sensitive name of a requirement. Another may require the use of an internally-generated serial number. A third may provide a persistent “handle” that can only be decoded by calling a proprietary API. To make use of any of these reference mechanisms, the referring tool must know both what kind of reference it is dealing with, and also to which tool it refers.

¹⁴ For a more rigorous and complete definition of REST, see Roy Fielding’s PhD dissertation, in particular Chapter 5, available online at http://www.ics.uci.edu/~fielding/pubs/dissertation/rest_arch_style.htm.

¹⁵ For example, if the user of a browser clicks on a link to a PDF document, the browser will ask Adobe Acrobat to handle the document that is returned from the server. If the link points to a document in OpenDocument Text (ODT) format, the browser will use whatever tool is available that supports ODT.

¹ In July 2008, Google announced that its web search surpassed one trillion unique URLs.

² Some commercial content providers have objected to a practice known as “deep linking” because it bypasses their ability to control a site visitor’s access path or to enforce subscription payments. However, this is an issue of business models, not of technology.

³ More information about the Linked Data concept can be found online at <http://linkeddata.org>

⁴ Ironically, the DVD standard was almost derailed before it even started by an eerily similar format war. The intervention of IBM, which rallied other computer vendors to agree to boycott both camps unless they compromised on a unified specification, is a little-known story in the success of DVD.

⁵ Such situations are so common that the IT industry needs a term for these “de facto” standards, as contrasted with “de jure” standards, which are standards established by an independent body and available equitably to any market participant.

⁶ Imagine how much the value of mobile phones would be diminished if, for example, customers of Sprint could only call other customers of Sprint, and customers of Verizon only other Verizon customers. Extraordinary as it now sounds, that is exactly how email worked in the early days of online access. Subscribers to one online service could only email customers on the same service. Even today, the instant messaging services provided by various web portal providers pursue this same “walled garden” approach.



Please Recycle