

What is the Semantic Web?

This is a book about a thing we call The Semantic Web. Just from the name, you can guess that it is related somehow to the familiar World Wide Web (WWW), and has something to do with Semantics. Semantics, in turn, has to do with understanding the nature of meaning, but even the word “semantics” has a number of meanings. In what sense are we using the word “Semantics”? And how can that be applied to the web?

This book is also about a Working Ontologist. That is, the aim of this book is not to motivate or pitch the Semantic Web, but to provide the tools that someone who plans to work with the Semantic Web will need. Or perhaps more accurately; the World Wide Web Consortium (W3C) has provided these tools in the forms of standard Semantic Web languages, complete with abstract syntax, model-based semantics, reference implementations, test cases, etc. But these are like a craftsman’s tools; in the hands of a novice, they can produce clumsy, ugly, barely functional output. But in the hands of a skilled craftsman, they can be used to produce works of utility, beauty, and durability. It is our aim in this book to describe the craft of building Semantic Web systems. We go beyond coverage of the fundamental tools to show how they can be used together to create Semantic Models, sometimes called *Ontologies* that are understandable, useful, durable, and perhaps even beautiful.

1.1 What’s a Web?

The idea of a web of information was once a technical idea accessible only to a highly trained elite of information professionals; IT administrators, librarians, information architects and the like. Since the widespread adoption of the WWW, it is

now common to expect just about anyone to be familiar with the idea of a web of information, shared across the world. Contributions to this web come from every source, and any topic under the sun is covered.

Essential to the notion of the web is the idea of an open community; anyone can contribute their ideas to the whole, for anyone to see. It is this openness that has resulted in the astonishing comprehensiveness of topics covered by the web. An information web is an organic entity that grows from the interests and energy of the community that supports it. As such, it is a hodge-podge of different analyses, presentations and summaries of any topic that suits the fancy of anyone with the energy to publish a web page. Even as a hodge-podge, the web is pretty useful. Anyone with the patience and savvy to dig through it can find support for just about any inquiry that interests them. But the web often feels like it is “a mile wide, but an inch deep.” How can we build a more integrated, consistent, deep web experience?

1.2 *Smart web, Dumb web*

Suppose you hit the web page for a major national park, and you find there a list of hotels that have branches in the vicinity of the park. In that list you see that Mongotel, one of the well-known hotel chains, has a branch. Since you have a Mongotel rewards card, you decide to book your room there. So you click on the Mongotel web site, and search for the hotel there. To your surprise, you can't find a Mongotel branch at the national park. What is going on here? “That's so dumb,” you tell you browsing friends, “if they list Mongotel on the national park website, shouldn't they list the national park on Mongotel's web site?”

Suppose you are planning on attending a conference in a far-off city. The conference web site lists the venue where the sessions will take place. You go to the website of your preferred hotel chain, and find a few hotels in the same vicinity. “Which hotel in my chain is nearest to the conference?” you wonder. “And just how far off is it?” There is no shortage of websites who can compute these distances, once you give them the addresses of the venue and your own hotel. So you spend some time copying and pasting the addresses from one page to the next, and noting the distances. “That’s dumb,” you think to yourself, “Why should I be the one to copy this information from one page to another. Why do I have to be the one to copy-and-paste all this information into a single map?”

Suppose you are investigating our solar system, and hit a comprehensive web site about objects in the system; stars (well, there’s just one of those), planets, moons, asteroids and comets are all described there. Each object has its own web page, with photos and essential information (mass, albedo, distance from the sun, shape, size, what object it revolves around, period of rotation, period of revolution, etc.) . At the head of the page is the object category, e.g., planet, moon, asteroid or comet. Another page includes interesting lists of objects; the moons of Jupiter, the named objects in the asteroid belt, the planets that revolve around the sun. This last page has the nine familiar planets, each linked to its own data page. One day, you read in the newspaper that the International Astronomical Union (IAU) has decided that the object known as Pluto, which up until 2006 was considered to be a planet, should be considered a member of a new category called “dwarf planet”! You rush to the Pluto page, and see that indeed, the update has been made – Pluto is listed as a Dwarf Planet! But when you go back to the

“Solar Planets” page, you still all see nine original planets listed under the heading “Planet” – Pluto is still there. “That’s dumb,” you think to yourself, “that they didn’t update the web pages consistently!”

What do these examples have in common? Each of them has an apparent representation of data, whose presentation to the end user (the person operating the web browser) seems “dumb.” What do we mean by “dumb”? In this case, “dumb” means inconsistent, out of synch, and disconnected. What would it take to make the web experience seem ‘smarter’? Do we need smarter applications, or a smarter web infrastructure?

1.3 Smart Web Applications

The web is full of intelligent applications, with new innovations coming every day. Ideas that once seemed futuristic are now commonplace; search engines make matches that seem deep and intuitive; commerce sites make smart recommendations personalized in uncanny ways to your own purchasing patterns. Mapping sites include detailed information about world geography, and can plan routes and measure distances. The sky is the limit for the technologies a web site can draw upon; every information technology under the sun can be used in a web site, and many of them are. New sites with new capabilities come on the scene on a regular basis.

But what is the role of the web infrastructure in making these applications ‘smart’? It is tempting to make the infrastructure of the web smart enough to encompass all of these technologies and more. The smarter the infrastructure, the smarter the web’s performance, right? But it isn’t practical, or even possible, for the web infrastructure to provide specific support for all, or even any, of the technologies that we might want to

use on the web. Smart behavior in the web comes from smart applications on the web, not from the infrastructure.

So what role does the infrastructure play in making the web smart? Is there a role at all? We have smart applications on the web – why are we even talking about enhancing the web infrastructure to make a smarter web, if the smarts aren't in the infrastructure?

The reason we are improving the web infrastructure is to allow smart applications to perform to their potential. Even the most insightful and intelligent application is only as smart as the data that is available to it. Inconsistent or contradictory input will still result in confusing, disconnected, “dumb” results, even from very smart applications. The challenge for the design of the Semantic Web is not to make a web infrastructure that is as smart as possible; it is to make an infrastructure that is most appropriate to the job of integrating information on the web. The Semantic Web doesn't make data smart – smart data isn't what the Semantic Web needs. The Semantic Web just needs to get the right data to the right place, so that the smart applications can do their work. So the question to ask is not “How can we make the web infrastructure smarter?” but rather, “What can the web infrastructure provide to improve the consistency and availability of web data?”

1.4 A connected web is a smarter web

Even in the face of intelligent applications, disconnected data results in dumb behavior. But the web data doesn't itself have to be smart; that's the job of the applications. So what can we realistically and productively expect from the data in our web applications? In a nutshell, we want data that doesn't surprise us with inconsistencies that make us want to say, “This page is dumb.” We don't need a smart web infrastructure – we just need a web infrastructure that lets us connect data to smart

web applications so that the whole web experience is enhanced. The web seems smarter, because smart applications can get the data they need.

In the example of the hotels in the national park, we'd like there to be coordination between the two web pages, so that an update to the location of hotels would be reflected in the list of hotels at any particular location. We'd like the two sources to stay synchronized. Then we won't be surprised at confusing and inconsistent conclusions drawn from information taken from different pages of the same site.

In the mapping example, we'd like the data from the conference website, and the data from the hotels website, to be automatically understandable to the mapping website. It shouldn't take interpretation by a human user to move information from one site to the other. The mapping website already has the smarts it needs to find shortest routes (taking into account details like toll roads and one-way streets), and to estimate the time required to make the trip. But it can only do that, if it knows the correct starting and end points.

We'd like the astronomy website to update consistently; if we state that Pluto is no longer a planet, the list of planets should reflect that fact as well. This is the sort of behavior that gives a reader confidence that what they are reading reflects the state of knowledge reported in the website, regardless of how they read it.

None of these things are beyond the reach of current information technology. In fact, it is not uncommon for programmers and system architects, when they first learn of the Semantic Web, to exclaim proudly, "I implemented something very like that for a project I did a few years back. We used . . .," and they go on to explain how they used some conventional, established technology such as relational databases, XML stores or object stores to make their data more connected and consistent. But what is it that these

developers are building? What is it about managing data this way that made it worth their while to create a whole subsystem on top of their base technology to deal with it? And where are these projects, two or more years later? When those same developers are asked whether they would rather have built a flexible, distributed, connected data model support system themselves, or to have used a standard one that someone else optimized and supported, they unanimously chose the latter. Infrastructure is something that one would rather buy than build.

1.5 Semantic Data

In the Mongotel example, there is a list of hotels at the national park, and another list of locations for hotels. The fact that these lists are intended to represent the presence of a hotel at a certain location is not explicit anywhere; this makes it difficult to maintain consistency between the two representations. In the example of the conference venue, the address appears only as text typeset on a page so that human beings can interpret it as an address. There is no explicit representation of the notion of an address, or the parts that make up an address. In the case of the astronomy web page, there is no explicit representation of the status of an object as a planet. In all of these cases, the data describes the presentation of information, rather than describing the entities in the world.

Could it be any other way? Can an application organize its data so that it provides an integrated description of objects in the world and their relations, rather than their presentation? The answer to this is “yes,” and indeed it is common good practice in web site design to work this way. There are a number of well-known approaches.

One common way to make web applications more integrated is to back them with a relational database, and generate the web pages from queries run against that database.

Updates to the site are made by updating the contents of the database. All web pages that require information about a particular data record will change when that record changes, without any further action required by the web maintainer. The database holds information about the entities themselves, while the relationship between one page and another (presentation) is encoded in the different queries. Consider the case of the National Parks and Hotel. If these pages were backed by the same database, the National Park page could be built on the query “Fetch all hotels with location=National Park”, and the hotel page could be built on the query, “Fetch all hotels from chain=Mongotel”. If Mongotel has a location at the National Park, it will appear on both pages; otherwise, it won’t appear at all. Both pages will be consistent. The difficulty in the example given is that it is organizationally very unlikely that there could be a single database driving both of these pages, since one of them is published and maintained by the National Park Service, while the other is managed by the Mongotel chain.

The astronomy case is very similar to the hotel case, in that the same information (about the classification of various astronomical bodies) is accessed from two different places, ensuring consistency of information even in the face of diverse presentation. It differs in that it is more likely that an astronomy club or university department might maintain a database with all the currently known information about the solar system.

In these cases, the web applications can behave more robustly by adding an organizing query into the web application, to mediate between a single view of the data and the presentation. The data itself isn’t any less dumb than it was before, but at least it is centralized, and the application or the web pages can be made to organize the data in a way that is more consistent for the user to view. It is the web page or application that

behaves smarter, not the data. While this approach is useful for supporting data consistency, it doesn't help much with the conference mapping example.

Another approach to making web applications a bit smarter is to write program code in a general-purpose language (e.g., C, Perl, Java, Lisp, Python or XSLT) that keeps data from different places up to date. In the Hotel example, such a program would update the National Park web page whenever a change is made to a corresponding hotel page. A similar solution would allow the planet example to be more consistent. Code for this purpose is often organized in a relational database application in the form of *stored procedures*; in XML applications it can be affected using a transformational language like XSLT. These solutions are more cumbersome to implement, since they require special-purpose code to be written for each linkage of data, but have the advantage over a centralized database that they do not require all the publishers of the data to agree on and share a single data source. Furthermore, such approaches could provide a solution to the conference mapping problem, by transforming data from one source to another. Just as in the query/presentation solution, this solution does not make the data any less dumb, it puts an informed infrastructure around the data, whose job it is to keep the various data sources consistent.

The common trend in these solutions is to move away from having the presentation of the data (for human eyes) be the primary representation of the data; that is, they move from having a website be a collection of pages, to having a website be a collection of data, from which web page presentations are generated. The application focuses not on the presentation, but on the subjects of the presentation. It is in this sense

that these applications are semantic applications; they explicitly represent the relationships that underlie the application, and generate presentations as needed.

1.6 A Distributed Web of Data

The Semantic Web takes this idea one step further; applying it to the web as a whole. The current web infrastructure supports a distributed network of web pages, which can refer to one another with global links called Uniform Resource Locators (URLs). As we have seen, sophisticated web sites replace this structure locally with a database or XML backend that ensures consistency within that page.

The main idea of the Semantic Web is to support the idea of a distributed web at the level of the data, rather than the presentation. Instead of having one web page point to another, one data item can point to another, using global references called Uniform Resource Identifiers (URIs). The web infrastructure provides a data model whereby information about a single entity can be distributed over the web. This distribution allows the Mongotel example and the conference hotel example to work like the astronomy example, even though the information is distributed over websites controlled by more than one organization. The single, coherent data model for the application is not held inside one application, but rather is part of the web infrastructure. When Mongotel publishes information about its hotels and their locations, it doesn't just publish a human-readable presentation of this information, but instead it publishes into the web infrastructure a distributable, machine-readable description of the data. The data model that the Semantic Web infrastructure uses to represent this distributed web of data is called the Resource Description Framework, or RDF, and is the topic of Chapter 3XXX.

This single, distributed model of information is the contribution that the Semantic Web infrastructure brings to making data less dumb. Just as is the case with data-backed web applications, the semantic web infrastructure allows the data to drive the presentation, so that various web pages (presentations) can provide views into a consistent body of information. In this way, the semantic web helps data not be so dumb.

1.7 Features of a Semantic Web

The World Wide Web was the result of a radical new way of thinking about sharing information. These ideas seem familiar now, as the web itself has become pervasive. But this radical new way of thinking has even more profound ramifications when it is applied to a web of data like the Semantic Web. These ramifications have driven many of the design decisions for the Semantic Web Standards, and have a strong influence on the craft of producing quality semantic web applications.

1.7.1 Give me a voice . . .

On the World Wide Web, publication is by and large in the hands of the content producer. Anyone can build a web page, and say on it what they please. A wide range of opinions on any topic can be found; it is up to the reader to come to a conclusion about what to believe. The web is the ultimate example of the warning, *caveat emptor*. This feature of the web is so instrumental in its character that we give it a name; we call it the *AAA Slogan*: “Anyone can say Anything about Any topic.”

In a web of documents, the AAA slogan means that anyone can write a page saying whatever they please, and publish it to the web infrastructure. In the case of the semantic web, it means that our data infrastructure has to allow any individual to express

a piece of data about some entity in a way that can be combined with information from other sources. This requirement sets some of the foundation for the design of RDF.

It also means that information is not managed as is usual for a large, corporate data center, where one database administrator rules with an iron hand over any addition or modification to the database. A distributed web of data, in contrast, is an organic system, with contributions coming from all sources. It was this freedom of expression on the document web that allowed it to take off as a bottom-up, grassroots phenomenon.

1.7.2 . . . so I may speak!

In the early days of the document web, it was common for skeptics, hearing for the first time about the possibilities of a world-wide distributed web full of hyperlinked pages on every topic, to ask, “But who is going to create all that content? Someone has to write those web pages!”

To the surprise of those skeptics, and even to many proponents of the web, the answer to this question was that *everyone* would provide the content. Once the web infrastructure was in place, so that Anyone could say Anything about Any topic, people came out of the woodwork to do just that. Soon every topic under the sun had a web page, either official or unofficial. It turns out, a lot of people had something to say, and they were willing to put some work into saying it.

The document web grew because of a virtuous cycle that is called the *network effect*. In a network of contributors like the web, the infrastructure made it possible for anyone to publish, but what made it *desirable* for them to do so? At one point in the web, when web browsers were a novelty, there was not much incentive to put a page on this new thing called “the web;” after all, who was going to read it? Why do I want to

communicate to them? Just as it isn't very useful to be the first kid on the block to have a fax machine (who do you exchange faxes with?), it wasn't very interesting to be the first kid with a web server.

But as a few people did have web servers, and a few more got web browsers, it became more attractive to have both web pages and web browsers. Content providers found a larger audience for their work; content consumers found more content to browse. As this trend continued, it became more and more attractive, and more people joined in, on both sides. This is the basis of the network effect; the more people who are playing now, the more attractive it is for new people to start playing.

A good deal of the information that will populate the semantic web is already available on the web, typically in the form of tables, spreadsheets or databases. Who will do the work of converting this data to RDF for distributed access? In the earliest days of the Semantic Web, there was little incentive to do so. As more and more data is available in RDF form, it becomes more useful to write applications that utilize this distributed data. The Semantic Web has been designed to benefit from the same network effect that drove the document web.

1.7.3 What about the Round-worlders?

The network effect has already proven to be an effective and empowering way to muster the effort needed to create a massive information network like the World Wide Web; in fact, it is the only method that has actually succeeded in creating such a structure. The AAA slogan enables the network effect that made the rapid growth of the web possible. But what are some of the ramifications of such an open system? What does the AAA slogan imply for the content of an organically grown web?

In order for the network effect to take hold, we have to be prepared to cope with a wide range of variance in the information in the web. Sometimes the differences will be minor details in an otherwise agreed-upon area; at other times, differences may be essential disagreements that drive political and cultural discourse in our society. This phenomenon is apparent in the document web today; for just about any topic, it is possible to find web pages that express widely differing opinions about that topic. The ability to disagree, and at various levels, is an essential part of human discourse, and a key aspect of the web that makes it successful. Some people might want to put forth a very odd opinion on any topic; someone might even want to postulate that the world is round, while others say that it is flat. The infrastructure of the web must allow both of these (contradictory) opinions to have equal availability and access.

There are a number of ways in which two speakers on the web may disagree; we will illustrate each of them with the example of the status of Pluto as a planet:

- They may fundamentally disagree on some topic. While the IAU has changed its definition of *Planet* in such a way that Pluto is no longer included, it is not necessarily the case that every astronomy club or even national body agrees with this categorization. Many astrologers, in particular, who have a vested interest in considering Pluto to be a planet, have decided to continue to consider Pluto as a planet. In such cases, different sources will simply disagree.
- Someone might want to intentionally deceive. Someone who markets posters, models, or other works that depict nine planets has a good reason

to delay reporting the result from the IAU, and even in spreading uncertainty about the state of affairs.

- Someone might simply be mistaken. Web sites are built and maintained eventually by human beings, and are subject to human error. Some web site might erroneously list Pluto as a planet, or indeed, might even erroneously fail to list one of the eight “non-dwarf” planets as a planet.
- Some information may be out of date. There are a number of displays around the world of scale models of the solar system, in which the status of the planets is literally carved in stone; these will continue to list Pluto as a planet until such time as there is funding to carve a new description for the ninth object. Websites are not carved in stone, but it does take effort to update them; not everyone will rush to accomplish this.

While some of the reasons for disagreement might be, well, disagreeable (wouldn't it be nice if we could stop people from lying!), in practice there isn't any way to tell them apart. The infrastructure of the web has to be able to cope with the fact that information on the web will disagree from time to time, and that this is not a temporary condition. It is in the very nature of the web that there be variations and disagreement.

1.7.4 To each his own

How can the web infrastructure support this sort of variation of opinion? That is, how can two people say different things, about the same topic? There are two approaches to this issue. First, we have to talk a bit about how one can make any statement at all in a web context.

The IAU can make a statement in plain English about Pluto, e.g., “Pluto is a dwarf planet.” But such a statement is fraught with all the ambiguities and contextual dependencies inherent in natural language. We think we know what “Pluto” refers to, but how about “dwarf planet”? Is there any possibility that someone might disagree on what a “dwarf planet” is? How can we even discuss such things?

The first requirement for making statements on a global web is to have a global way of identifying the entities we are talking about. We need to be able to refer to “The notion of Pluto as used by the IAU” and “the notion of Pluto as used by the American Federation of Astrologers” if we even want to be able to discuss whether the two organizations are referring to the same thing by these names.

In addition to Pluto, another object was also classified as a “dwarf planet”. This object is sometimes known as UB313, and sometimes known by the name “Xena”. How can we say that the object known to the IAU as “UB313” is the same object that its discoverer Michael Brown calls “Xena?”

One way to do this would be to have a global arbiter of names decide on how to refer to the object. Then Brown and the IAU can both refer to that “official” name, and say that they use a private “nickname” for it. Of course, the IAU itself is a good candidate for such a body, but the process to name the object has already taken over two years. Coming up with good, agreed-upon global names is not always easy business.

But in the absence of such an agreement, different web authors will select different URIs for the same real-world resource. Brown’s Xena is IAU’s UB313. When information from these different sources is brought together in the distributed network of data, the web infrastructure has no way of knowing that these need to be treated as the

same entity. The flip side of this is that we cannot assume that just because two URIs are distinct, that they refer to distinct resources. This feature of the semantic web is called the *Non-unique naming assumption*; that is, we have to assume (until told otherwise) that some web resource might be referred to using different names by different people.

1.7.5 There's always one more

In a distributed network of information, as a rule we cannot assume at any time that we have seen all the information in the network, or even that we know everything that has been asserted about one single topic. This is evident in the history of Pluto and UB313. For many years, it was sufficient to say that a “Planet” was defined as any object orbiting the sun of a particular size. Given the information available during that time, it was easy to say that there were nine planets around the Sun. But the new information about UB313 changed that; if a planet is defined to be any body that orbits the sun of a particular size, then UB313 had to be considered a planet, too. Careful speakers in the late 20th century of course spoke of the “known” planets, since they were aware that another planet was not only possible but even suspected (the so-called “Planet X”, which stood in for the unknown but suspected planet for many years).

The same situation holds for the semantic web. Not only might new information be discovered at any time (as is the case in solar-system astronomy), but also, because of the networked nature of the web, at any one time a particular server that holds some unique information might be unavailable. For this reason, on the Semantic Web we can rarely conclude things like “there are nine planets”, since we don't know what new information might come to light.

In general, this aspect of a web has subtle but profound impact on how we draw conclusions from information we have. It forces us to consider the web as an *Open World*, and to treat it using the *Open World Assumption*. An open world in this sense is one in which we must assume at any time that new information could come to light, and we may draw no conclusions that rely on assuming that the information available at any one point is all the information available.

For many applications, the open world assumption makes no difference; if we draw a map of all the Mongotel hotels in Boston, we get a map of all the ones we know of at the time. The fact that Mongotel might have more hotels in Boston (or might open a new one) does not invalidate the fact that it has the ones it already lists. In fact, for a great deal of semantic web applications, we can ignore the open world assumption, and simply understand that a semantic application, like any other web page, is simply reporting on the information it was able to access at one time.

The openness of the web only becomes an issue when we want to draw conclusions based on distributed data. If we want to place Boston in the list of cities that are not served by Mongotel (e.g., as part of a market study of new places to target Mongotels), then we cannot assume that just because we haven't found a Mongotel listing in Boston, that no such hotel exists.

As we shall see in the following chapters, the semantic web includes features that correspond to all the ways of working with open worlds that we have seen in the real world. We can draw conclusions about missing Mongotels, if we say that some list is a comprehensive list of all Mongotels. We can have an anonymous "Planet X" stand in for an unknown but anticipated entity. These techniques allow us to cope with the open

world assumption in the semantic web, just as they do in the open world of human knowledge.

1.8 Summary

The aspects of the web we have outlined here – the AAA slogan, the network effect, non-unique naming and the open world assumption – already hold for the document web. As a result, the web today is something of an unruly place, with a wide variety of different sources, organization and styles of information. Effective and creative use of search engines is something of a craft; efforts to make order from this include community efforts like social bookmarking and community encyclopedias to automated methods like statistical correlations and fuzzy similarity matches.

For the semantic web, which operates at the finer level of individual statements about data, the situation is even wilder. With a human in the loop, contradictions and inconsistencies in the document web can be dealt with by the process of human observation and application of common sense. With a machine combining information, how do we bring any order to the chaos? How can one have any confidence in the information we merge from multiple sources? If the document web is unruly, then surely the semantic web is a jungle; a rich mass of interconnected information, without any roadmap, index or guidance. How can such a mess become something useful? That is the challenge that faces the working ontologist. Their medium is the distributed web of data; their tools are the semantic web languages, RDF, RDFS and OWL, Their craft is to make sensible, usable and durable information resources from this medium. We call that craft *modeling*, and it is the centerpiece of this book.

On the cover of this book, we see a system of ducts, conduits and channels with water coursing through them. If we think of the water as the data on the web, the channels are the model. If not for the model, the water would not flow in any systematic way; there would simply be a vast, undistinguished expanse of water. Without the water, the ducts and conduits would have no dynamism; they have no moving parts in and of themselves. Put the two together, and we have a dynamic system. The water flows in an orderly fashion, defined by the structure of the ducts and channels. This is the role that a model plays in the Semantic Web. Without the model, there is an undifferentiated mass of data; there is no way to tell what data can or should interact with other data. The model itself has no significance without data to describe. Put the two together, however, and you have a dynamic web of information, where data flows from one point to another in a principled, systematic fashion. This is the vision of the Semantic Web – an organized World Wide system where information flows from one place to another in a smooth but orderly way.

1.9 Fundamental Concepts

The following fundamental concepts were introduced in this chapter:

The *AAA Slogan* – Anyone can say Anything about Any topic. One of the basic tenets of the web in general, and the semantic web in particular.

Open World/Closed World – a consequence of the AAA slogan is that there could always be something new that someone will say; this means that we must assume that there is always more information that could be known.

Non-Unique naming – since the speakers on the web won't necessarily coordinate their naming efforts, the same entity could be known by more than one name.

The *Network Effect* – the property of a web that makes it grow organically. The value of joining in increases with the number of people who have joined, resulting in a virtuous cycle of participation.