Website Location: Cyberspace vs. Geographic Space

Gareth Sansom
Department of Justice Canada
Draft: April 3, 2008

The problem of what it means to speak of a Website's **location** has changed over the past decade and a half. If we are to understand how this problem has changed we must begin with a model that is simple but with sufficient explanatory power to be useful. Figure 1 represents the building blocks of location.

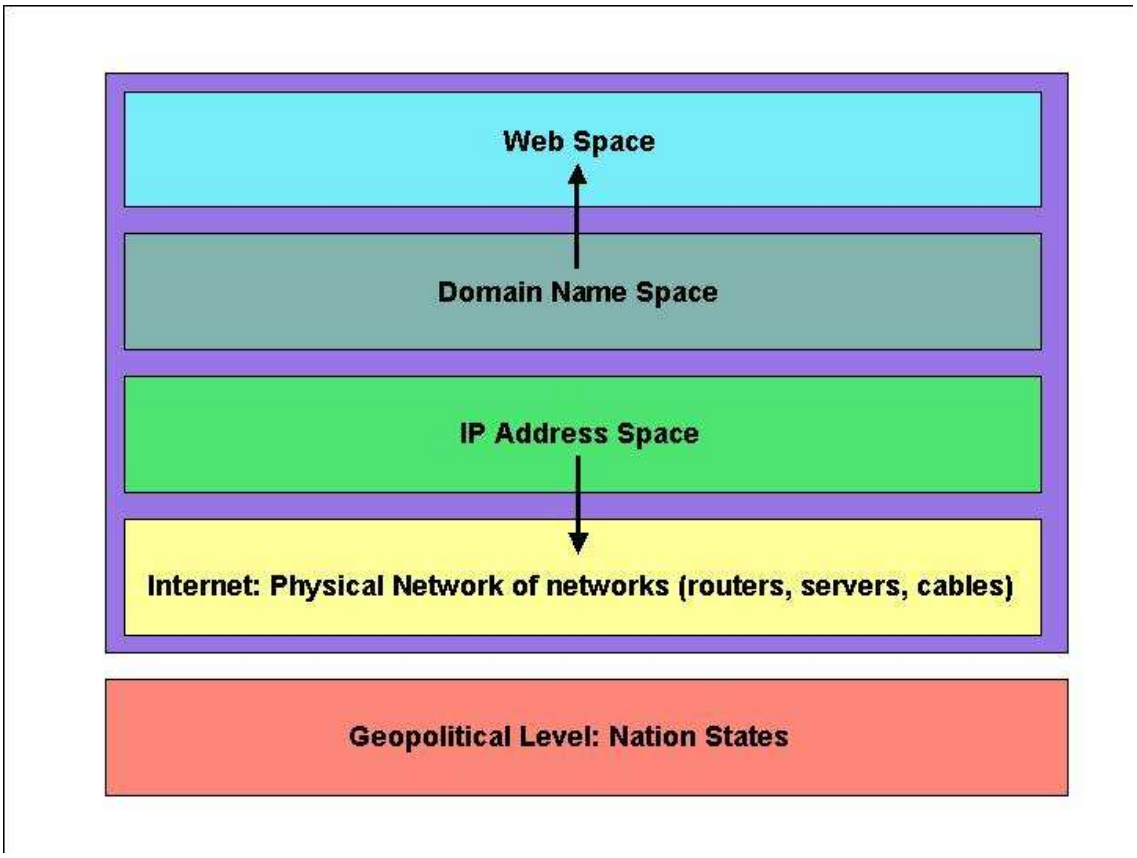

Figure 1. Different Spaces: The Building Blocks of Location

Identifying the location of a Website is often relevant because one is really interested in the Website's content. As a consumer one may wish to access the content. As a state one may wish to block access to content deemed to be illegal or one may wish that a Website be taken down. All these wishes have different probabilities of success depending on the circumstances and the measures taken. Content, once digitized, can take a broad range of forms (text, images, video, etc.) on a broad range of storage media (paper, CD-ROMs or DVDs, hard drives, thumb drives, etc.). In the age of digital reproduction, although

individual instances of content ("tokens") can be destroyed, it is often difficult to destroy *content as such* because it can be easily and accurately copied, creating new tokens of the same content. To the extent that content can exist offline, even if actions have been taken to remove it from the Internet, it can be quickly replaced from any location where it exists (including mobile devices such as laptops, PDAs, DVDs) to one or more new locations (servers) on the Internet. Content that was formerly fragile or ephemeral can now have persistence over time by virtue of its reproducibility.

If everything is static – static content, static hosting and static addressing – then location is stable. In a static world, location in cyberspace can be made to coincide with location in geopolitical space. A static situation of this sort would have been an accurate description of an ftp archive sitting on the Internet twenty-five years ago. Webspace, however, has inherent fluidity: it is dynamic by nature. "WebSpace" is the space formed by hypertext or hypermedia linkages, it forms a complex web of linked information resources. We should remind ourselves that a **stable** location does not mean **single** location. With hypermedia, the content of a single, static Web page can come from more than one location: the Web page is hosted on a certain server but the images could be hosted on a different server (in another country, for example): what is perceived as one document by the end-user can actually come from **multiple locations**. Assume one Web page with hyperlinks that remained stable, assume also the Websites to which those links point (domain names and pages) all continue to exist, **however** those Websites themselves could have changed from one month to the next, whether in terms of their content or in terms of the computers hosting that content (hence not only the machines changed but location could have changed, security could have changed, ownership could have changed). In this sense, WebSpace is dynamic at the **global level**.

If everything becomes dynamic, then location becomes more problematic. If circumstances shift from static content, static hosting and static addressing to dynamic content, dynamic hosting and dynamic addressing, then the relationship between the layers or strata in Figure 1 cease to be static and become dynamic. Location then becomes a placeholder with multiple possible values. This forces the interested party to decide which location among this plurality is relevant and for what purpose.

Let us touch on dynamic content briefly. If hypermedia documents are dynamic, the documents themselves can change based on different factors, such as user input, user history or even advertisement revenue:

–the text, pictures or video in the document can change
–the fields in a form can change
–the content inside the portion of the page containing an advertisement can change and different advertisements can be served

Such content changes can occur at the user end (**client-side scripting**) or at the provider end (**server-side scripting**). If we look to dynamic content on the server side, we should conceive the Web page as an "enframing document". Let us assume that the "enframing" document remains on the serving computer and that computer has a fixed location: only in this respect is there one location that is stable. The content that is piped into the various areas of the enframing document, however, can change: that

content can be served by different computers at different times (possibly changing minute to minute) – these computers could be situated in different locations around the world and could be owned by different companies or different persons than whoever is responsible for the enframing document. In normal business cases, who can serve what is established by contract and structured by the underlying technology. The end-user perceives a single document on their screen which changes just like watching a program on a television channel BUT unlike a television channel, each part of the screen could be coming simultaneously from different computers around the world: connect to the same page an hour later and the content might have changed as well as the locations from which that content is being served. Server-side dynamic content opens the possibility that the content from a single Web page comes from multiple locations that change on the fly from one moment to the next.

The relationship between the domain name and the IP address can become dynamic: meaning the domain name remains the same but the IP address associated with it changes. Dynamic addressing permits dynamic hosting. One changes the IP address, because one wants to change the host server – this could be for good business reasons (load balancing) or for criminal reasons (evasion). In the business world, dynamic DNS is provided by large scale DNS hosting servicee which retain the current addresses in a database and provide a "client" program to the user which will send an update to the service whenever the server's IP address has changed. So the domain name remains the same but the computer hosting the content changes:

–the server hosting the content could be in France;
–24 hours later the server hosting the content could be in Germany;
–24 hours later the server hosting the content could be in Japan

This was the standard technique a few years ago for child pornography pay-per-view Websites and for botnet command-and-control servers.

If we distinguish the server hosting the content from computers serving the content (as in a “fast flux service network”) then the switchover can be even faster:

– the computer serving the content could be in France;
– 3 minutes later the computer serving the content could be in Germany;
–3 minutes later the computer serving the content could be in Japan

This is a new technique used by botnets in the past year or two: the serving computers are compromised computers scattered around the world. A “double flux service network” uses one fast flux network to serve the content and a second fast flux network to manage the dynamic DNS addressing. If what is being served was a Web page containing illegal content, to the end-user it would appear that the Website and its address (badstuff.com) are stable. Yet in these cases there are multiple (hundreds or thousands) of computers that are not legally owned by the perpetrators; these multiple computers are in multiple locations in multiple jurisdictions**.**