What is *Content-Based Routing?*

A Semandex Information Note

Introduction

Everyone is familiar with the internet model of networking, even if few of its users understand the details of IP naming and addressing. However, the range of services that are available on the internet are implicitly determined by the capabilities of the infrastructure on top of which the global internet is implemented. In order to understand how these services can be extended, or enhanced, we need to explain a little of the current model.

Basic IP

The basic IP service is the delivery of a packet of data from one IP address to another IP address. IP itself does not care about the contents of that packet. It just delivers it. On top of the IP layer, the transport protocols, TCP or UDP, add additional functionality to this basic service, such as reliability and byte-streaming.

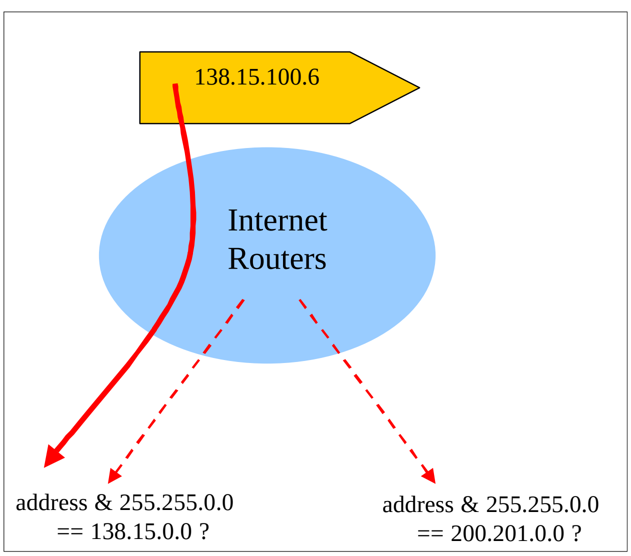

Figure 1. Addressing in the IP Internet

Of course, each machine in the world does not know how to route packets to all of the 2^{32} (in IPv4) other machines in the universe. IP was designed to be a scalable service, each machine knows about its own neighborhood, and it knows where to send packets that are going outside its neighborhood. Once inside the network, IP routers send packets successively closer to their destination. IP routing relies on the 32-bit IP address having a structure in which machines that are close to each other, in networking terms, have addresses that are also close to each other. This is the reason behind the familiar class-based routing model and netmask values (which define the IP addresses that are close enough to be in the same neighborhood) shown in figure 1.

Neither do machines (or humans) know the IP addresses for the endpoints of all potential connections. Instead, the IP addresses are looked up as needed, from a machine name which is intended to be remembered by humans.

This system has worked very well for many decades, and will continue to serve the needs of point-to-point connections where the endpoint names are known or easily determined. One area in which the original IP model has been severely strained is in broadcast distribution services, where a single source wishes to send identical data streams to multiple consumers. The multiple-unicast (N separate streams) approach is very inefficient in bandwidth, and cannot scale to more than a few tens of users. The alternative, IP multicast, solves the bandwidth and scalability issue by only duplicating packets as needed usually when the same packet must be sent out on two different ports from the same router. There are three main problems with the multicast approach:

1. Multicast uses special IP addresses, which limits the number of multicast channels that can be in simultaneous operation;

2. Users need to locate and join a multicast group using a special protocol;

3. If there is a break in the multicast tree, or a packet is lost, all the downstream consumers will request a retransmission of that packet. This leads to the so-called ACK-implosion, which remains a topic of research within the IP community.

Currently, the applications of multicast are limited to video or other real-time services, but the duplication of transmitted information brought about by the World-Wide-Web may force the viability of multicast to be re-evaluated.

Enter the Web

Perhaps one of the most surprising aspects in the evolution of the internet is that the Web 'revolution' has not required any changes to the basic IP model. A Web URL (Uniform Resource Location) is essentially an extension of a machine name to access a particular resource (a service, a file, a web-page etc.) on that machine. From the network view, that is all there is to the Web, a set of point-to-point connections for accessing data.

What has made the Web so much more that just data, as far as its users perceive it, is the standardized representation of pages on the Web that permits a (fairly) consistent presentation of information across a range of devices and software platforms. Added to this is the built-in concept of a clickable link that allows users to remain mostly unaware of IP, names, addresses, protocols and services. More than the network, the Web is this standard representation, html, a derivation of the Standard Generalized Markup Language (SGML), and the semantics of the <href> tag that it defines. Understanding, or accepting, this concept is the key to realizing how Content-based routing can enter the field, and take it over.

One problem that the Web, in itself, does not solve, is how to find information. The fact remains that, unless you know the URL for a resource, you cannot access that resource. In the early days, with few sites, this was less of an issue, but in today's Web, locating information only based on known URLs would be impossible. The Web solution is, of course, the Search Engine, as shown in figure 2. A Web search engine is basically a repository of URLs with some indication of the contents of that URL. Users submit a query to this URL database, and receive a set of records for the URLs whose content is believed to match the query. The user can then visit each or all of the indicated resources until the required information is located.

The disadvantages of the current search engines are well documented, even though it is clear that the current Web would not function without them. The main issues are coverage and accuracy. A Search Engine cannot return a URL that is not in its database. If a search engine visits only 25% of the Web, then it might be plausibly assumed that there are three times as many relevant pages out there as

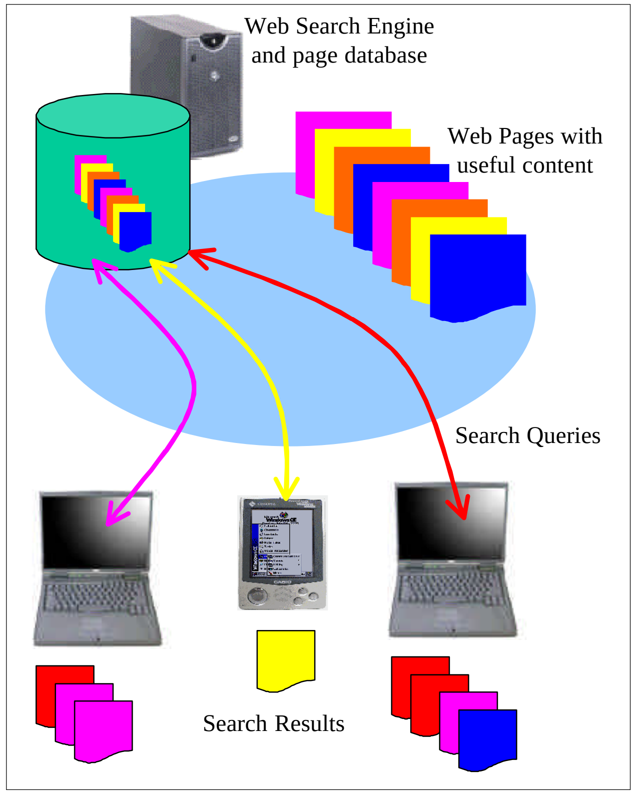

Figure 2. The Web and Search Engines

that engine can return. However, as most search users quickly realize, quantity is a very poor substitute for quality.

Relevance measures are familiar to the Information Retrieval community, and generally express the ratio of wanted answers to unwanted answers. For most queries in search of information, the only measure is how quickly can that information be found in the list, i.e. what is the rank of the first one, or possibly two, relevant results. Getting the answer on hit 2 or 3 is considered superior to getting it from hit 42 or 43. In this type of query, the number of 'false positives' at the top of the list needs to be minimized. Less common, but also useful, is the "exhaustive search" seeking to find all relevant information "leaving no stone unturned" for example in seeking out prior publications for patent filings. Here, an omitted URL may be costly. The need then is to minimize the number of false negatives in the list.

The Portal Approach

Global search engines are rarely the right vehicle for this second type of search, in which the query operates in a restricted domain, and where more specific tagging of content is needed than the generalized keyword crawlers of the Net. Such requirements are usually better served by domain-specific 'portals,' which can monitor information from multiple sources, correlate the results, and deliver only relevant output to individual subscribers. As this implies, portals differ from search engines in that they require individual subscriptions to their service. Aside from the commercial nature of these operations, the need for subscriptions is clear. In order to filter the incoming (often real-time) information, the portal needs to know whether or not it is relevant to one or more users. This, in turn, requires that each user supply an 'interest profile' to the service. The portal stores the profiles and, based on the users' preferences, builds up a list of relevant documents which the user can later retrieve.

For many applications in this area, a portal is close to the ideal solution. The subscriber is guaranteed a reliable, fairly complete set of documents addressing a specific interest or set of interests. Indeed, the portal suffers from very few drawbacks:

1. Portals are usually domain-specific, in order to maximize the accuracy of their classifications. This means that a single consumer may require subscriptions to multiple portals to cover a wider range of interests.
2. The portal provider knows the interest profiles of all its subscribers. In a commercial environment this is highly sensitive information. Exposing the profiles of corporate executives is one reason why some companies decline to use such services.

The portal concept, personalized delivery of relevant, timely, information provides the service model for its evolution towards Content-based delivery.

In an ideal world, each user would be their own portal, able to specify an interest profile and receive only the information that matches that profile. However, this is at odds with the current point-to-point nature of the Internet. To achieve this goal, two components must be in place:

1. Information content in documents must be expressed in a way that can be used to match interest profiles of users
2. Elements within the network must be capable of delivering documents to users based on the match between the content of the document and the interest of the user.

Taken together, these two services provide the essential elements for Content-based routing.

XML - the key to content

XML is to content what html is to layout. Also derived from SGML, the eXtensible Markup Language (XML) provides a standardized way of describing the content of a document. As its name implies, rather than trying to define a single standard covering all possible contents, XML is a framework within which domain-specific descriptions or schemas can be defined, to extend the basic capabilities. Using a domain-specific schema allows meaning to be attached to words or phrases in a document. For example "Fleet" within a financial schema may be tagged as the name of an institution, whereas "Fleet" in a naval context would represent a grouping of vessels. This allows the context of "The Fleet Bank opened a new branch" and "The Fleet skirted Dogger Bank" to be disambiguated (a vital first stage in reducing the number of false positives for a query on "Fleet" AND "Bank" in the financial world).

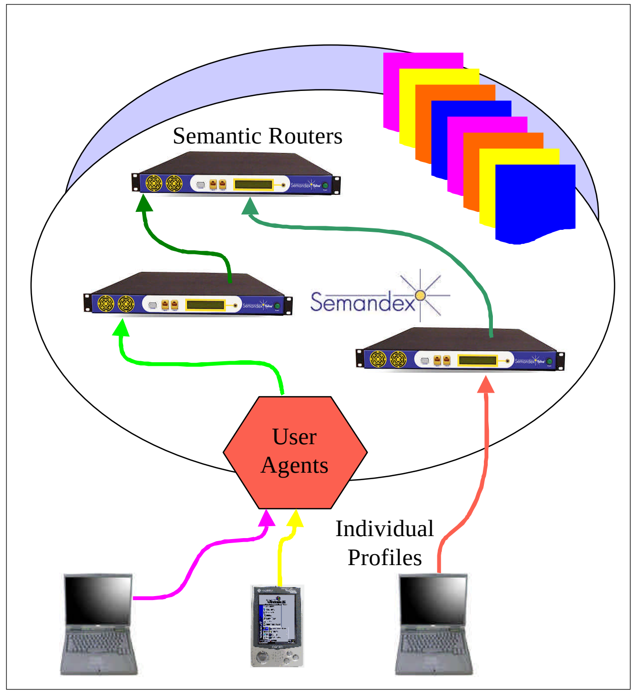

Figure 3. Interest Profile Information Flow

XML does more than simply tag words. XML schemas can also tag fields and other structured content, for example to mark a value as a price, in dollars, or a stock-quantity remaining, or a discount as a percentage etc. Given data tagged in this way, it becomes possible to contemplate'asking the net to find "a red Corvette car for sale with under 50,000 miles priced between $10,000 and $15,000 within 25 miles of my house," without getting any responses from the Navy about upgrading their Corvette-class destroyers.

Semandex Netlink™ - the key to routing

Documents still need to reach their targets, and there is nothing in XML that will map a document onto a set of IP addresses for the consumers that should receive it. Indeed, in a content-based routing system, the producer of the information has no idea who, or where, the potential consumers are. What is required is a parallel architecture to the IP layer in which data packets are routed towards interested consumers within the network. Similar to IP multicast, if the same information is desired by users on different ports of a router, the packet must be duplicated to ensure each recipient gets a copy. The only essential difference is the nature of the routing table, instead of matching an IP address with an entry in a routing table to find the next hop towards the destination, the content-based router matches the description of the packet's contents with the interests of the users in order to find the next hop system. The interest profiles must first be stored on the routers, to build up these tables, as shown in Figure 3.

Of course, in traditional IP systems, route matching is done using addresses and netmasks in a relatively straightforward algorithm. In content-based routing, the XML description of a packet must be compared with the interest profile (also in XML) for each output port of a router. This comparison is considerably more complicated that the basic "AND-then-compare" IP system of Figure 1.

Semandex Networks Inc has developed just such a routing system, called Netlink. Based on a scalable line of content-based routers, and front- and back-end support software, Netlink runs as a Semantic Network on top of an existing IP infrastructure, just like the current Web. Packets are routed between end-users and Netlink routers, and between routers, using traditional IP addressing, but inside the box all the routing decisions are based on the content descriptors of the packets as shown in Figure 4.

Netlink is a truly scalable system, potentially growing to the size of the Internet, based on our proprietary routing technology. It is also a totally distributed system, each router handles the users in its neighborhood, passing packets towards the core of the network and delivering information to its community. Because there is no central portal, not only can the system grow quickly, but there is no central repository of interest profiles. Individual user profiles are known to the neighborhood router, true, but upstream routers see only the bulk, or aggregated, profile for all downstream users. In many cases, the local routers are further protected by company firewalls (see our "Semantic Firewalling" document in this same series).

The Future is now. Content-based routing and the Semantic Network.

We believe that content-based routing, properly understood, holds the promise for the next generation of the internet, beyond the Web. Just as the Web has established itself as a community of machines riding on the IP internet, so the Semantic Network will provide the solutions to the information location and dissemination needs for multiple information communities. Of course the Web and URLs will still exist, but we see the decline of search engines, to be replaced by "Semantic Seek" packets, being routed in real time in search of precisely-determined information, and triggering its delivery to the connected users.

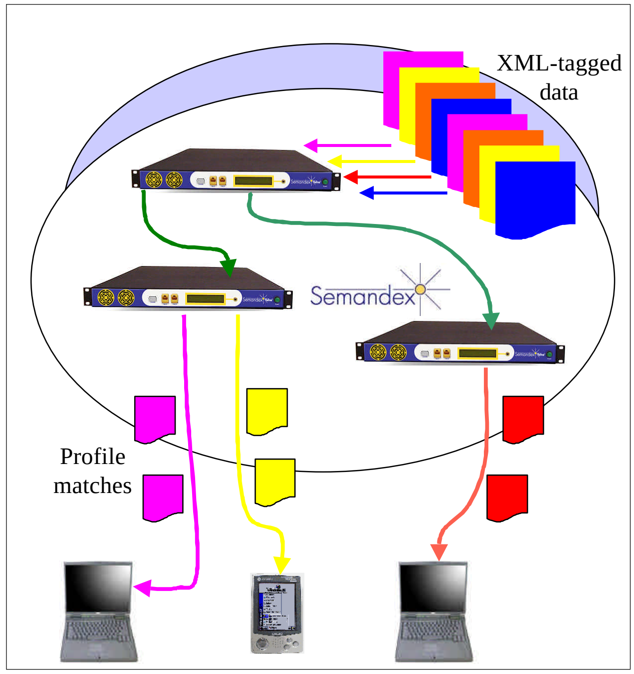

Figure 4. Data Flow from Producer to Consumers

Already this network is a reality, systems based on our Netlink technology have been used to disseminate information in demonstrator and trial systems. As more and more systems adopt XML as their data representation standard, the need for portals will diminish and ultimately all content delivery will be handled by semantically-aware components in the infrastructure, truly building "The Network That Knows."

For Further Information

Semandex Networks Inc.
201 Washington Road
Princeton, NJ 08540-6449 USA
www.semandex.net

Tel: 1-866-7SEMANDEX (1-866-773-6263)
Fax: 609-514-4061

Sales/Business Development: sales@semandex.net
Support: support@semandex.net