# Content Aware Networking in the Internet: Issues and Challenges

Barani Subbiah
3Com Corporation
barani_subbiah@3com.com

Zartash Afzal Uzmi
Stanford University
zartash@leland.stanford.edu

*Abstract- Content aware networking is a new paradigm in which the network is intelligent to understand the type of content request, where to find it, and how to deliver it in an efficient way. We discuss, why the existing IP-smart paradigm does not meet the needs of todays content rich Internet. We describe content aware networking issues that include content storage across distributed servers, content based routing, content delivery, content distribution and content based resource reservation. Content delivery is often associated with multicast and we present various multicast techniques by which a content can be delivered efficiently. The emerging end-to-end IP Security (IPSec) poses some unique challenge to this trend and we present few solutions to overcome the problem.*

## 1. Introduction

The success of todays Internet is largely to the vast amount of contents available at no cost to users. Internet traffic measurements have shown that content access is the dominant service in todays Internet [1]. As the number of users in the Internet increase, so as the number and diversity of contents. However, todays networking protocols and devices do not meet the needs of the content related services. Current services on the Internet are limited to those in which a connection is established based on the IP addresses of the machines. The dominant routing protocols in the Internet such as Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) are capable of routing packets based on IP addresses [2]. However, these protocols have no knowledge of which server (IP address) is suitable for a particular content.

In the present Internet architecture, it turns out that being IP-smart only is not being smart enough. In addition to being IP-Smart, being content smart is quite beneficial in various circumstances. To appreciate the benefits of content aware networking, we will examine a scenario where a portal site has more than one content server managed through a Domain Name Server (DNS). Current DNS implementations return IP addresses, of multiple servers with same domain name, in a round robin scheme. DNS has no knowledge if these servers have different processing capacity and/or load. This scheme is clearly inefficient as compared to the one in which a DNS returns the IP address of the server which is either less congested or one that is more appropriate for that particular content type.

Content delivery is an another important issue in which a content aware network is more beneficial. The network can route different type of contents among different routes and reserve resources without the user or application level signaling. For example, if the network is aware of voice as the content, then it can route it along a prioritized path as well as reserve resources accordingly. The emergence of end-to-end IPSec while accessing content in the Internet is posing some unique challenges to this trend. Many of the areas mentioned above have issues in common. We will address these in order to develop a complete understanding of content aware networking and challenges in realizing it in the Internet. Furthermore, the previous work for realizing content related services is scattered, and different areas have been addressed independently, in a manner that it is difficult to draw a relationship among them.

## 2. Content Service Model

A content service model in a content driven Internet consists of Content customer, Content Service Provider (CSP) or Content Broker (CB) and Content Provider (CP). These three entities are inter-connected through service level agreements. CSP offer services by which a customer is able to access the content. The most common services of a CSP include, locating a content, searching static contents to build a dynamic content, charging information, and content negotiation on behalf of a customer with a CSP. In current Internet, portal services such as Yahoo, America On-Line (AOL), and Infoseek can be considered as CSPs. We believe that CSPs will remain as an integral part of the content driven Internet. The single most reason is due to the vast presence of contents in the Internet and CSP is the only entity that can help a customer to access content services efficiently and economically. CP in this model refer to the actual creator or owner of the content. In addition to the above three entities, we see that Internet Connectivity Provider (ICP) also plays a role in this model. It is quite possible that a ICP, CSP and CP are a single entity such as AOL.

## 3. Content Storage and Content Request Distribution

The enormous increase in the number of users and contents in the global Internet has led to a cluster based server architecture, as shown in Figure 1. In this model, several back-end servers are used to store and deliver contents and a single front end server (distributor) is responsible for receiving and distributing content requests. Distributing the incoming requests to the back-end servers in a round-robin way is quite inefficient considering that some of the content is requested most of the time and most of the content is requested some of the time. Some commercial products are already available which are slightly smarter than the round-robin schemes. Cisco's LocalDirector product distributes the incoming requests based on the load on the back-end servers [3]. This scheme performs significantly better than the round-robin scheme, however it does not give consideration to the requested content type. Another product, Connect-Control, by Check Point does perform load balancing in a similar fashion with additional freedom of allowing user selected load balancing algorithm but this product also does not give any consideration to the content type and locality of previously requested content [4].

Locality Aware Request Distribution (LARD) scheme [5] considers the locality of the requested content and the load distribution on back-end server at the same time and, therefore, gives a much higher throughput than any of the schemes used in currently available commercial products. While LARD scheme focuses on the static content, HACC Architecture [6] also takes dynamic content into account. In HACC Architecture, the front-end server is termed as smart router whose function is to identify one of the back-end servers that should satisfy an incoming request and then routing the incoming request to that particular server.

One potential problem with a cluster architecture is that the front-end requires as much networking resources as collectively required by all nodes, which may render the front-end to be network bottleneck [7]. For example, TCP handoff procedure which limits the number of states that can be maintained at the front-end. A new scheme, as shown in Figure 2, is being proposed to distribute the intelligence of the front-end router to back end servers [7]. In this model, an incoming content request can be received by any server in the cluster. Each server will determine which server is suitable in terms of load and content type match so that it can hand off the request. This requires that all servers communicate with each other through a protocol to learn the status of other servers. The advantage of this scheme is that there is no bottleneck front end as every server share the incoming load. However, a new protocol is required between servers which might increase traffic load as well as modification necessary at the servers to accommodate the new functions.

We see that the present solutions of load balancing are a short term approach. In order to be a true content aware networking model, each content request should be analyzed for its content type. Servers should be arranged based on content type such as audio and video server (multimedia servers). Also, servers that specialized for static and dynamic contents would further improve the content services.

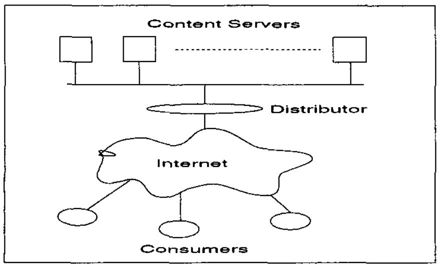


Figure 1: Centralized model

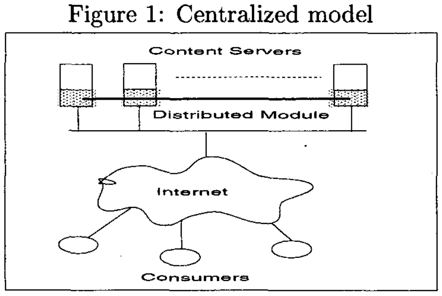


Figure 2: Distributed model

## 4. Content delivery and Content distribution

Even though content delivery and content distribution refer serving contents to customers, we observed a small difference between them. It is becoming common to use content delivery when a content is served directly to customers. However, content distribution refers to distributing content from one server to multiple servers (mirror sites) located at different locations. It is becoming a common practice that content is pushed or stored close to the customers, which means that content servers are distributed at the edges of network. There is much debate between content caching and distributed content servers in terms of efficiency, which is out of scope for this discussion.

### 4.1 Multicast content delivery

One of the problems with multicast content delivery is that how to serve a static content such as streaming video to a group of users whose requests are received at different times at a content server. Few schemes have been proposed in earlier literatures and we will analyze them in detail. The use of time slot multicast to deliver contents to a group of users was proposed in [8]. This mechanism can be applied to both reliable and unreliable multicast to deliver popular documents (web pages) to a group of users who requested the same page at the same time or within a small interval of time. The time is divided into small chunks called as time slots. All the users whose requests, for the same content, are received in the same time slot would be referred to as simultaneous users. Since the requested content is assumed to be very popular, it is very likely that several requests would be received in a time slot. Once a group of users is identified then the server can create a multicast group and instruct them to join a new multicast group. This new group is served separately by the server with that content. It can be seen that next time slot will result in a new multicast group served by the same server with the same content. It is quite possible that there will be large number of multicast groups present at any given time for the same content. Once a content is served, the corresponding multicast group is terminated. The advantage of this method is that the servers do not worry about packet order and clients can start the play sequence immediately since they always receive the content from the beginning.

By the definition of time slot, it is clear that longer time slots result in longer wait time for the users and shorter time slots take less advantage of the multicasting mechanism by having fewer requests in shorter time slots. Therefore, an optimum time slot size must be selected based on these two non-orthogonal parameters, user wait time and multicast efficiency. Since there is usually an upper limit for user wait time, maximum size of a time slot is fixed. The multicast delivery mechanism would be beneficial if the minimum expected number of requests exceed a certain threshold in a time slot. The minimum number of simultaneous users can be computed in a time slot [8], therefore it can be estimated whether multicasting would be beneficial or not. It is easy to see that the use of multicast becomes more likely as the requested content becomes more popular.

In addition to the above scheme, a cyclic best effort multicast scheme was also investigated [9]. This method is based on a continuous multicast without needing any requests from individual users. The use of cyclic best effort multicast was proposed for very popular web sites, typically the top-level page of a very popular site. As the name suggests, the cyclic best effort mechanism involves transmitting the same content over and over until all (or most) of the requests have been fulfilled with a high probability. One complete transmission of the whole content is defined as one cycle. The required number of cycles to satisfy all requests with a given probability can be computed by using a discrete time markov chain model. Similarly, same model can be used to compute the expected number of satisfied receiver. While cyclic multicast can be used effectively for downloadable images and video, and even for live broadcast where the transmission is needed just for one cycle.

### 4.2 Multicast content distribution

As we pointed out earlier, content distribution is between content servers. Recently, reliable multicast has been proposed to distribute contents especially between content servers with in the realm of a web portal. As the content servers are distributed and moved closer to the edges of a network in order to reduce the latency, we see that reliable multicast is being used to push the content from one server to all other servers. In addition, it can also be used to deliver any reliable content to a group of users such as stock information. It is expected that content distribution, as the killer application, will facilitate a wide spread deployment of multicast in the Internet.

## 5. Content aware routing

Content routing pertains to directing the request to the most appropriate server with intelligence closer to the client [10]. The first part of content based routing is that which server(s) has the content. Once a server is identified then the node closest to user can find out how to route the request. A simple scenario of content based routing is shown in Figure 3, which consists of a client, a network and a single server distributed as S1 and S2. The server distribution may be non-overlapping or it may be complete duplication (often referred to as mirroring). It should be noted that S1 and S2 are two different machines with two different IP addresses and are located at two different geographical locations. Each of these may represent a cluster of servers in which there is a front-end server and a number of back-end servers. Current implementation of DNS points to a group of IP addresses for a single fully qualified hostname, and a name resolution request from the client to the DNS is returned with the IP address of one of the servers in a round-robin fashion. This scheme is obviously inefficient for two reasons. First, the requested content may not lie on the server whose IP address is returned by the DNS to the client for non-duplicated content. Second, the connection to the server, whose IP address is returned by the DNS, might be much slower than some other server in case of fully duplicated (mirrored) content.

A commercial product that attempts to address the content routing problem is Ciscos DistributedDirector [11]. While the DistributedDirector suggests a better IP address than suggested by the round-robin DNS scheme by giving consideration to network delays and other performance parameters, it still does not give any consideration to the routing based on the requested content type. The problem is, therefore, to design an interface to the DNS, which hopefully remains transparent to the client and returns the IP address of the most appropriate server that needs to be contacted. For such an interface to work

0-7803-7097-1/01/$10.00 ©2001 IEEE

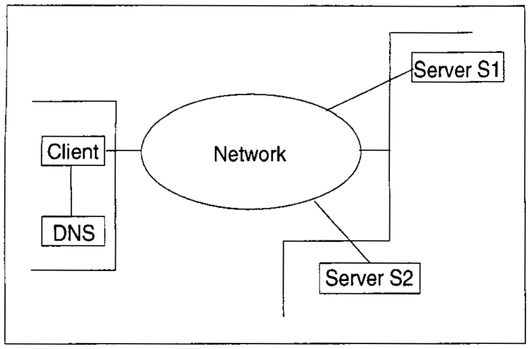


Figure 3: A general network connection model

properly, an accurate extraction of metadata from the content itself is also of primary importance.

## 6. Location aware content services

Location aware content services might answer the questions like: What and where are the restaurants near me? Currently available services do answer questions like this but they require that the user type in their current location before using any location aware queries. One prime example of this type of service is the www.sidewalk.com. If the location of the user can be ascertained by use of, for example, Global Positioning System (GPS) receivers, then this information can be piggybacked with the location aware query [12]. The use of clients location is also helpful for content routing, such that the interface to a content aware DNS respond with the IP address of one of the distributed servers which is closest to the client. Here the word closest is used in the network sense and includes several parameters that characterize the speed of a connection, including but not limited to number of hops and end-to-end delay. In fact, location aware services are more useful when a CSP want to push contents to customers. This allows CSP to provide up to date information about a customer's environment such as a sale in a nearby shopping Mall or traffic report as the user enters a traffic zone.

One of the problems with providing location aware content services is due to lack of association of IP addresses with the physical location. Two alternate solutions can be proposed to satisfy location-based queries. First, each client knows its physical location, perhaps in terms of longitude and latitude, by making use of GPS or some other device and this location information is transmitted to the CSP. CSP then processes the location information along with the requested content and contacts the content provider (CP) to retrieve the desired content and finally delivers it to the client, or redirects the client to the desired content.

Following scenarios are foreseen to provide location aware content services. For non-dialup connections, the location of default gateway is usually fixed and can be transmitted as approximate location of the client, which is assumed to be close to the default gateway. This certainly requires quite a bit of cooperation among already deployed protocols. For the dialup connections, on the other hand, exact location can always be known even if a GPS receiver is not utilized. This is possible by making use of 911 like services in which a telephone number is uniquely indexed against a physical address, and the phone number is also ascertained by making use of caller ID. This is an elegant solution to realize location based services and the only challenge is its implementation on top of existing protocols without modifying the clients. For mobile clients, this is not a problem since existing techniques allow a cellular operator to track a mobile with in a serving area. This information can then be transmitted along with the content request to a content server.

## 7. Security and content aware networking

IP Security define ways to encrypt data between peer entities thus protecting integrity and privacy of the data [13]. The use of IPSec is becoming popular due to widespread deployment of Virtual Private Networks (VPN) based services as well as increasing concern over customer privacy in public Internet. If IPSec is used between peer nodes, intermediate networking nodes no

0-7803-7097-1/01/$10.00 ©2001 IEEE

longer have access to the content carried inside an IP packet. In a sense, IPSec defeats the purpose of content aware networking. This is one of the reasons why content aware networking may not be feasible in the core of the Internet. However we can solve security issues in a CSP domain such as portal sites.

As we described earlier, content distribution involves Front End Servers (FES) and Back End servers (BES) where FES handles all the incoming content requests. Our studies have shown that if CSPs want to use content aware networking techniques with IPSec, it is advisable that IPSec is terminated at the FES. Once FES decrypts and analyzes the content, it becomes much easier to use content aware networking techniques to route, reserve and charge accordingly. Another way to solve this problem is to identify the content type in the IP header, e.g extension to IPv6 header. For example, voice can be identified as one type or even multiple types based on the codecs and video can be another type. However, it will limit the number of content aware networking ideas that can be used.

## 8. Conclusions

Content aware networking is becoming an integral part of content rich Internet, as networking moves from routing (layer3) to applications and services (layer7). It is becoming critical that the network understands the dominant content related services that it is being used for. We explained, why the current IP smart paradigm does not meet the needs of the emerging content rich Internet. We believe that CSPs are an integral part of the content service model and we expect this model to continue as long as diversity and loosely managed content structure are retained in the Internet. Content delivery and distribution are often associated with multicast and we see that these applications will usher widespread multicast deployment in the Internet. Content routing has scalability issue in a global internet, however it can easily be applied in a web portal realm. It is certain that a content aware Internet is beneficial to both users and network/content providers as the number and diversity of the content increases in the Internet.

## References

[1] S. lawrence and C.L. Giles: Accessibility of Information on the web: Nature, Vol 400, July 1999.

[2] C. Huitema: Routing in the Internet, Prentice Hall PTR, ISBN - 0130226475

[3] Cisco LocalDirector, Available at http://www.cisco.com

[4] Check Point Connect Control, Available at http://www.checkpoint.com

[5] V.S. Pai et al: Locality Aware Request Distribution in cluster based network servers, SIGPLAN Notices, Vol.33, no.11, pp.205-16, 0326-1340

[6] X. Zhang, M. Barrientos, J. Bradley, and M. Seltzer: HACC: An architecture for cluster based web servers, 3rd USENIX Windows NT symposium

[7] L. Aversa and A. Bestavros: Load balancing a cluster of web servers, IEEE IPCCC2000 conference proceedings, Feb 2000.

[8] R.J. Clark and M.H. Ammer: Providing scalable web servers using multicast communication, Computer networks and ISDN systems, Elsevier, vol.29, no.7, pp.841-58, Aug. 1997.

[9] K.C Almeroth, M.H. Ammer and Z. Fei: Scalable delivery of web pages using cyclic best effort multicast, Proc. of IEEE infocom 98, pp. 1214-21, vol.3, 1998.

[10] M.A. Sheldon: Content routing: A scalable architecture for network based information discovery, Ph.D. Thesis, MIT.

[11] Ciscos Distributeddirector, Available at http://www.cisco.com

[12] T. Imielinski and J.C. Navas: GPS based geographic addressing routing and resource discovery, Communications of the ACM, vol. 42, no. 4, pp. 86-92, April 1999.

[13]Internet RFCs 2402-24011, Available at http://www.ietf.org/html.charters/ipsec-charter.html