

# Content Services Network: The Architecture and Protocols

Wei-Ying Ma  
Microsoft Research China  
wyma@microsoft.com

Bo Shen and Jack Brassil  
Hewlett-Packard Laboratories  
{boshen, jtb}@hpl.hp.com

## Abstract

Content delivery networks (CDNs) can be viewed as application-specific overlay networks that make web caching an infrastructure service accessible to any content provider. As the Internet continues to evolve with increasing diversity and heterogeneity, we see a growing demand for extending the capabilities of network intermediaries to provide additional services such as content adaptation, personalization, watermarking and location-aware data insertion.

A content services network (CSN) is proposed in this paper to make content transformation and processing an infrastructure service accessible to its subscribers. One can think of CSN as another layer of network infrastructure built around CDNs. This layer interacts collaboratively with user-agents, content servers, and other network intermediaries including ISPs' caching proxies and CDNs' surrogates in the content delivery process to provide value-added services. Furthermore, a CSN provides network resources that are used as a "service" distribution channel for value-added services providers to make their applications an infrastructure service.

To demonstrate the utility of our proposed CSN model, we describe the prototype implementation of a video segmentation and keyframe selection system as an infrastructure service that can be used by content providers or end users to enhance the way video is delivered over the Internet.

**Keywords:** Content delivery network, Internet proxy services, Internet Content Adaptation Protocol, Internet service delivery

## 1 Introduction

In the past two years we have seen a huge industrial effort devoted to the research and development of content delivery networks (CDNs). One can think of CDNs as an overlay network of a large number of caching proxies deployed at the network edge and used as a distribution channel to push content closer to the end user. CDNs handle many issues that used to be a technical burden to the content provider, such as scalability (service automatically scales up while the client requests increase), reliability (content is always available), and flash crowd control. CDNs enable content providers to avoid having to provision and manage their own networks and servers to handle the above issues.

In CDNs, caching is the primary intermediary service performed between clients and servers. That is, CDNs only replicates content and makes sure it is delivered reliably and promptly to end-users. They do not modify the content by adding value to it. As the Internet continues to evolve with increasing diversity and heterogeneity, we see a growing demand for extending the capabilities of network intermediaries to provide additional

services such as content adaptation, personalization, watermarking and location-aware data insertion. Due to the lack of network support, these problems are being ignored or resolved in an ad-hoc manner. For instance, to make content created for desktop PCs appear appropriately on handheld devices with smaller form factor, content providers typically re-author the content and create multiple versions for different devices. This process is not only time-consuming but also complicates content management.

To make content services an Internet infrastructure service, we propose a new architecture called "content services network (CSN)." One can think of content services network as another layer of network infrastructure (an overlaid network) built around CDNs. This layer interacts collaboratively with user-agents, content servers, and other network intermediaries including ISPs' caching proxies and CDNs' surrogates in the content delivery process to provide value-added services.

A CSN provides network resources that are used as a "service" distribution channel for value-added services providers to make their applications an infrastructure service. In many aspects, the benefit a CSN offered to value-added service providers are similar to what CDNs have offered to content providers. The services of a CSN can be subscribed and used by content provider, end user, ISP or even CDNs who pay for the service. For example, a content provider may request a content adaptation service from a CSN, and its content will be automatically adapted at the network edge before delivered to clients.

A brief comparison between CSN and CDNs is provided in Table 1. This comparison reflects a trend described by Katz [11], "the new Internet will be shaped by the ability to manage computation and storage deep inside the network, connected by application-specific overlay networks, all on behalf of end user applications. This is what we call services..." One can think of CDNs as application-specific overlay networks for storage and web caching. In comparison, CSN moves one step further and makes computation (processing and transcoding) an infrastructure service through the development of another overlay network around CDNs.

The technical issues a CSN needs to address are:

1. How to manage a large scale of computational resources and network services;
2. How to make the value-added services an infrastructure service that is accessible to its customers which could be content provider, end user, ISP or CDNs;
3. How to make a CSN interact with other existing network elements collaboratively and seamlessly so that the service model does not undermine the success of end-to-end nature of Internet client/server interactions.

	Content services network (CSN)	Content delivery networks (CDNs)
System overview	An overlay network of application proxies	An overlay network of caching proxies
Services	Processing	Storage and caching
Resource types	CPU and processor	Memory and disk
Protocol	Simple object access protocol (SOAP) and Internet service delivery protocol (ISDP)	Internet content adaptation protocol (iCAP)
Distribution channel for	Value-added services and applications	Web contents
Customers	Content providers, end users, ISPs, and CDNs	Content providers or ISPs

**Table 1.** Comparison of CSN and CDNs.

### 1.1 Related Works

The problem of adaptive content delivery in a heterogeneous network environment has been studied quite extensively in the past few years. Example works include Spyglass [14], ProxiNet [16], Intel QuickWeb [15], IBM Transcoding proxy [20], UC-Berkeley TranSend [12], Digestor [18], Mobeware [17], and Smart Client [21]. However, these works did not deal with the problem from the perspective of infrastructure service. Instead, they looked into the deployment of content adaptation technology at an origin server or a transcoding proxy.

In web caching and content distribution community, a number of Internet-Drafts have been proposed to address the problem of extending the functionality of a caching proxy for providing additional services that mediate, modify, and monitor object requests and responses. The IETF Working Group on Open Pluggable Edge Services (OPES) [23] and the Internet Content Adaptation Protocol (iCAP) [7] are the works closely related to our proposed CSN. We use some of the extended functions in their proposed future caching proxy as a foundation to build CSN.

Here we summarize the novel contributions of our work:

1. The idea of content services network (CSN) is proposed first time in the literature. We describe its system architecture and discuss two network protocols used by CSN to interact with other network elements for performing value-added services.
2. CSN provides an application framework in which SOAP [31] or iCAP operates, so we are able to investigate the issues such as control of policy. We propose a new Internet service delivery protocol (ISDP) which supplements these two protocols.
3. As an example of our CSN's service, we turn a video segmentation and keyframe selection system into an infrastructure service that can be used by content providers or end users to enhance the way video is delivered over the Internet.

## 2 System Architecture

The architecture of content services network (CSN) consists of the following three basic elements as shown in Figure 1. They are:

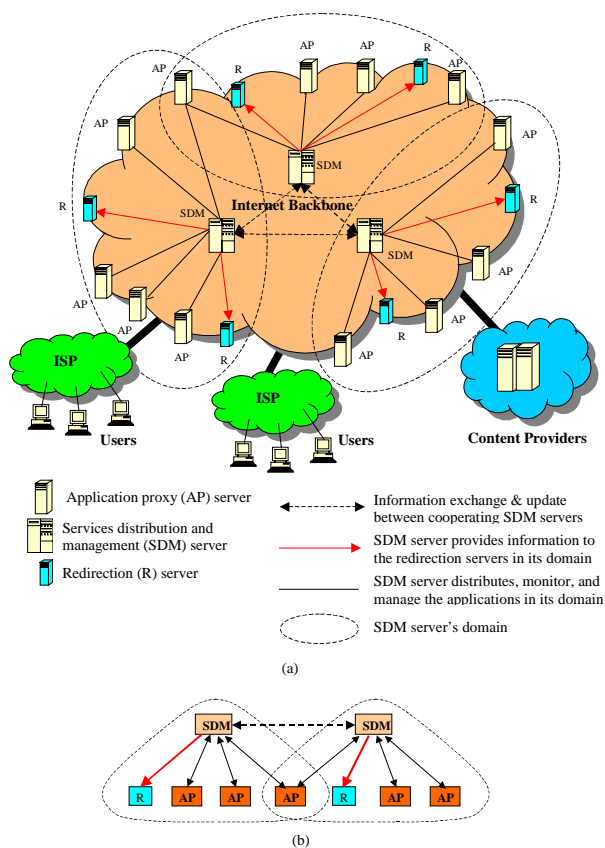
**Application proxy (AP) servers:** host the software of value-added services for content delivery. These application proxy servers are deployed at the edge of the Internet backbone, providing computational resources to process content on behalf of content providers (content servers) or end users (clients). If the service is provided to ISPs, it is considered performed on behalf of end-users. If the service is provided to CDNs who may resell it as a value-added service in addition to content delivery service, it is considered performed on behalf of content providers. Application proxy servers have a variety of operational modes depending on how the service is performed. They may behave like transcoding proxy servers, remote callout servers, or surrogate servers (with processing and special hosting capability). We will discuss the various service modes in each service mode in Section 4.

**Services distribution and management (SDM) servers:** are responsible for the following tasks: (a) they maintain static information about the service execution environment and the location of the applications, (b) SDM servers register and distribute the services (applications) provided by value-added services providers. They take into account a set of metric that measures the demand of services, the history of servers' load, demography of clients etc, and decide a set of application proxies to remotely install applications. (c) After initial service registration and distribution, SDM servers continue to monitor the performance of each service and dynamically adapt the scale of distribution and deployment of the service according to its demand. (d) SDM servers aggregate the information about usage pattern, availability and location of each deployed service, and then provide the information back to the redirection servers. The redirection servers use this information to perform service discovery, network and server load balancing and etc. (e) SDM server also provides management, accounting and billing functionalities to value-added service providers who use the CSN as service distribution channel. For example, value-added service providers may temporarily bring down their applications for maintenance or upgrade. They may want to know who have used their services and how much they have used so they can bill them. (f) Each SDM server is responsible for collecting information about its domain and periodically exchanges the information with other cooperating SDM servers. This information exchange can be triggered automatically if there is a change in the system. To reduce bandwidth consumption for information exchange, each SDM server may build a spanning tree that propagates the information efficiently to other SDM servers. Figure 1(b) shows the information exchange between two cooperating SDM servers.

**Redirection servers:** direct a service request to an application proxy server according to a number of attributes and measurements. Note that redirection servers receive information from SDM servers so they are able to perform services discovery and network and server load balancing

while directing service requests to an application proxy that is close to it in the network. The redirection servers are deployed at the edge of the network. When a service request is sent to a CSN, it is first directed to a redirection server that is close to it in the network. This redirection server then checks various attributes and measurements and further redirects the request to an application proxy server.

Note that CSN is developed as another layer of network infrastructure around CDNs. From the inside, this layer manages its network resources and applications (value-added services) to enhance the way content is distributed. From the outside, it defines interfaces and protocols to handle a variety of system interactions occurring in different types of service request and rendering. We will discuss the role of the application proxy server and its interaction with caching proxies in the next section. The protocol issues will be discussed in Section 4. The redirection problem is covered in Section 5.



**Figure 1.** The system overview of content services network (CSN)

### 3 Open Pluggable Edge Services (OPES)

In today's Internet, caching is the primary intermediary service performed between clients and servers to improve performance for web page access. As Internet continues to evolve with increasing diversity and heterogeneity, there is a need to extend the capabilities of network intermediaries for performing additional services that cannot be provided directly on clients or servers. The OPES has been formed to define protocols and APIs for a broad set of services that facilitate efficient delivery of complex content and services related to content [6].

Note that the most basic building component in a CSN is the application proxy server. In order to make CSN an open framework and architecture for distributing and executing value-added services developed by third-party companies, the application proxy servers should implement the protocols and APIs defined by OPES.

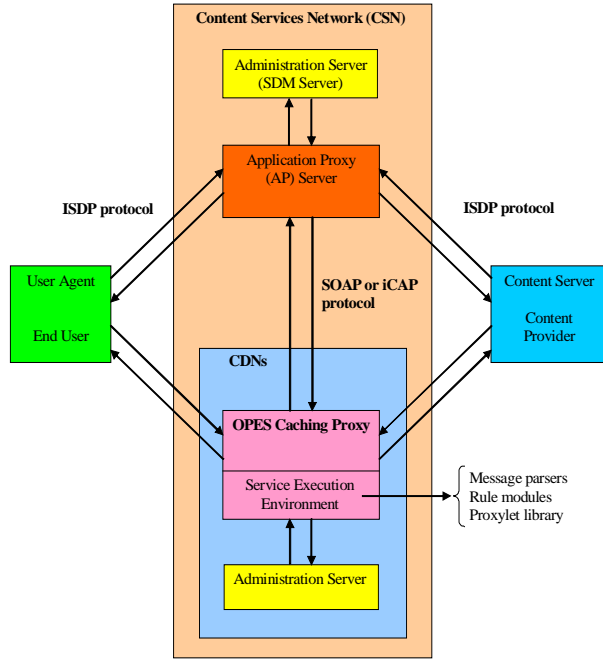
The current architecture of OPES supports two types of services; one is co-located with the transit intermediary (i.e. caching proxy) and the other is located on other cooperating servers (i.e. remote callout servers). For co-located services, OPES is defining the "service environment caching proxy" (referred as "*OPES proxy*" hereafter) which extends the functionality of a traditional caching proxy, making it capable of executing additional services that mediate, modify, and monitor object requests and responses. For remote callout services, OPES may use protocols like SOAP [31] or iCAP [7] that handle transport of objects to cooperating servers and back.

The OPES proxy is equipped with message parsers, rule modules and a proxylet library for general services. Its services can be extended through the downloading of rule modules and proxylets from user agents and content servers. While messages flow through an OPES proxy, they are parsed and matched against the rules specified in the rule modules. The matching of a rule causes a service to be performed on the message. The execution of the service may be performed by a local proxylet or by a remote callout server via a transport protocol.

Although one can think of the application proxies in a CSN as the remote callout servers in OPES, CSN and OPES have very different view of the world with respect to how content services may be distributed, subscribed and rendered. This difference reflects how the service path is formed and handled. OPES is very caching-proxy centric, and it assumes caching proxy continuously to be the center of the system, coordinating all the participating network elements. This setting limits its capability to address the services (and business models) that may require a more flexible arrangement of service path and system interaction among participating parties. In contrast, CSN shifts the focus to the "application proxy." CSN considers many possible service scenarios and defines a variety of service modes (i.e. service paths) to characterize the way the service may be subscribed and rendered. We will discuss the various service modes in Section 4.

Figure 2 shows the relationship between the OPES framework and the CSN architecture. Note that the application proxy is able to communicate directly with content servers and user agents to position itself at the right injection point in the service path to provide services. Keep in mind that the rendering of the service (processing) does not necessarily happen right after the service subscription stage. This is particularly true for post-distribution services such as personalization and insertion of regional data that are rendered at the network edge (more detail in Section 4). A typical scenario is that content servers or user agents may subscribe to the desired services from one application proxy server and receive the services rendered from the other application proxy server at a later time. In this case, a certificate containing the description of the subscribed services and authentication information is first given to content servers or user agents. This certificate is distributed with the target object or the client request, instructing other network elements (e.g. OPES caching proxy) to perform functions (e.g. running a

proxylet) or take actions (e.g. transporting objects to an application proxy via SOAP or iCAP protocol in order to receive value-added services).



**Figure 2.** The relationship between OPES proxies and the CSN architecture.

#### 4 Internet Service Delivery Protocol

Figure 2 shows two network protocols used in CSN. The SOAP and iCAP protocols provide an interface for CDNs and any caching proxy that supports OPES to communicate with a CSN for service delivery. One can think of SOAP or iCAP as a protocol for executing a “remote procedure call” on HTTP messages for some sort of adaptation. Basically, the SOAP or iCAP client sends a request along with a target object (like POST method in HTTP/1.1) to the application proxy server for transformation or value-added services. An example iCAP URI might look like:

`icap://ms1.csn.hp.net/video_services/summarization_and_keyframe_selection`

The SOAP and iCAP protocols are used along with a certificate issued by CSN for providing content services at the edge of network. This certificate represents a service to be performed on behalf of content providers or end users. On the other hand, the Internet Service Delivery Protocol (ISDP) provides an interface for content servers and user agents to request and receive the services from CSN. This certificate is part of the ISDP protocol to enable post-distribution service.

The Internet Service Delivery Protocol (ISDP) defines the interface between the application proxy (AP) servers and CSN’s customers for service subscription and rendering. In our current design, ISDP has three different operation modes based on

whom the service is performed on behalf of and the nature of the service. Note that each service mode corresponds to a different service path and system interaction between participating network elements.

A service could be performed on behalf of end-users (clients) or content providers (content servers). If the service is provided to ISPs, it is considered performed on behalf of end-users. If the service is provided to CDNs who may resell it as a value-added service in addition to content delivery service, it is considered performed on behalf of content providers. For the service performed on behalf of content providers, it is further distinguished based on whether it is “pre-distribution” or “post-distribution.” As the terms indicate, one way to compare these services is to look at when and where the value-added processing is conducted.

Pre-distribution services correspond to the processing that does not need to differentiate, for example, end-user or client’s capability. These services are performed before content is distributed over the Internet, and the responsible AP server acts like a surrogate (reverse proxy) to the origin content server. Pre-distribution services are usually performed once for an object and the result is reused for all subsequent requests. Watermarking [2] and logo-insertion are typical services of this type.

Post-distribution services correspond to the processing that requires a special handling on the object for individual client request. These services are subscribed first and rendered later at the edge of network. Post-distribution services are performed after content is distributed over the Internet, and the responsible AP server acts like an interception proxy or a remote callout server and is usually located in the client’s proximity. Post-distribution services may take into account the situation of each request, such as its location, end-user’s preferences and device’s capability, and then adapt the response accordingly. Example services of this type include content adaptation, insertion of regional data, customization and personalization.

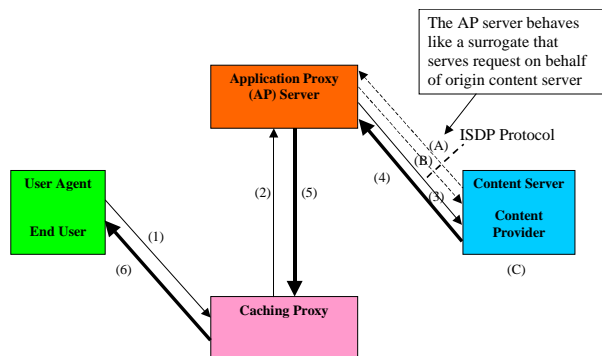
Table 2 shows the classification of a number of example proxy services described in [4] based on our service framework. We now describe how ISDP is operated in different service mode.

Example Proxy Services	Mode		
	I	II	III
Content adaptation for alternate web access devices		x	x
Insertion of Ad banners	x	x	
Insertion of regional data	x	x	
Language translation		x	x
Content adaptation for limited client bandwidth		x	x
Virus scanning			x
Content adaptation for alternate browser types		x	x
Adaptation of streaming media		x	x
Personalization and customization of web pages		x	x
Request filtering and content filtering			x
Watermarking	x		
Data hiding for enabling new applications	x		
Progressive multi-resolution image delivery (e.g., JPEG-2000)	x		

**Table 2.** The classification of example proxy services into the three service modes defined by a CSN. Mode I: pre-distribution service performed on behalf of content provider, Mode II: post-distribution service performed on behalf of content provider, Mode III: service performed on behalf of end user.

#### 4.1 Pre-distribution Service Performed on Behalf of Content Provider

Figure 3 depicts the service path for pre-distribution service performed on behalf of content provider. To set up a service path like this, the content server first sends an ISDP request to a CSN for a service for its object. This request is routed to a designated application proxy (AP) server through redirection. After receiving the request and knowing what service is needed, the designated AP server sends an ISDP response back to the content server. The ISDP response contains information such as digital signature for authentication and the IP address of the designated AP server for configuring the content server to use this AP server as its surrogate [1][5] for that object. Depending on the implementation, this step may involve rewriting the object's URL on content server to point directly to the designated AP server.



(A)	Content server sends an ISDP request to CSN to subscribe a service for its object. This request is routed to a designated AP server through redirection
(B)	The designated AP server sends an ISDP response back to the content server. The ISDP response contains information such as digital signature for authentication and the IP address of the designated AP server for configuring the content server to use this designated AP server as its surrogate for that object
(C)	The configuration may involve rewriting the object's URL on content server to point directly to the designated AP server. Note that the service relation between the AP server and the origin server is transparent to other network entities. That is, other network elements should send the requests for this object to the AP server as if it were the origin server.

(1)	A client sends a HTTP request for that object. The request goes to a caching proxy to see if there is a cache hit
(2)	Assuming no cache hit or the cached version has expired, the caching proxy forwards the request to the AP server because it thinks the AP server is the origin content server for the requested object
(3)	Assuming the object has not been value-added yet or the result of previous processing has expired, the AP server sends a request to fetch the object from the content server
(4)	The object is delivered to the AP server. The AP server run the subscribed processing on that object
(5)	The AP server sends the value-added object with expiration time to the caching proxy which may decide to cache it
(6)	The value-added object is delivered to the client

**Figure 3.** Pre-distribution service performed on behalf of content provider (content server).

In this mode, the designated AP server serves request for that object on behalf of the origin server, performing necessary processing to add value to the object before it is delivered over the Internet. Note that the service relationship between the AP server and the origin server is transparent to other network entities. That is, other network elements should send the requests for this object to the AP server as if it were the origin server. If CDN's surrogate proxies (for web caching and replication) are also used to accelerate content delivery for the origin server, they should talk to the designated AP server to get that object.

The processing for pre-distribution service is usually performed once and the result is cached for reuse. The designated AP server makes sure the result of processing remains fresh. It may fetch the object from origin server and run the processing again if the result has been expired or the object has been outdated.

The AP server performing this service mode is often equipped with a special server for hosting the value-added content as part of the service. For example, a web content provider may have a database of satellite images and aerial photographs that are typically large in size (in the order of 1000x1000 pixels). It is best rendering those images progressively and interactively, e.g. allowing users to zoom in a region of an image without downloading the whole image. The content provider can ask a CSN to provide the service that converts images into a special format, e.g. JPEG-2000, and then host them with progressive and interactive delivery capabilities.

Note that in this case, the communication between caching proxy and CSN's application proxy is a regular HTTP connection, and the caching proxy does not need to be "OPES-enabled."

#### 4.2 Post-distribution Service Performed on Behalf of Content Provider

Figure 4 depicts the system interaction occurred in the post-distribution service performed on behalf of content provider. Note that the service subscription and service delivery are conducted at different application proxy (AP) servers. In contrast to the previous mode, this type of service requires the collaboration of caching proxies at the edge of network to help render the service. That is, it needs the caching proxy to be "OPES-enabled." The CSN's AP server behaves like the remote callout server described in OPES.

To set up a service path like this, the content server first sends an ISDP request to CSN for a service for its object. After the redirection, one AP server responds the request by sending a certificate back to the content server. This certificate contains the description of the requested service and rules to perform it as well as the digital signatures necessary for performing authentication by the OPES proxy and the AP server which renders the service. This certificate is distributed together with the object. This certificate is used by the OPES caching proxies to determine when to perform a service, what kind of service to perform, and from where. An example ISDP certificate issued to a content provider might look like this:

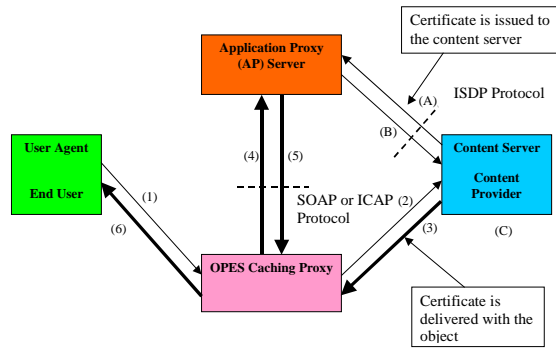
```
<subscriber = "content provider">
  <name>www.example.com</name>
<object>http://www.example.com/video/financial_result_Q2.m
pg</object>
```

```

<rule>
  <property name = "user-agent" matches = "HP Jornada |
Win CE">
</property>
</rule>
</subscriber>
<action>icap://ms1.csn.hp.com/video_services/summarization_a
nd_keyframe_selection</action>
<signature>AA%$#&!*@!G%N</signature>
</property>
</rule>
</subscriber>

```

In the above example, the content provider (e.g. [www.example.com](http://www.example.com)) has subscribed a service (e.g. icap://ms1.csn.hp.com/video\_services/summarization\_and\_keyframe\_selection) for its object (e.g. [http://www.example.com/video/financial\\_result\\_Q2.mpg](http://www.example.com/video/financial_result_Q2.mpg)). The certificate describes the requested service, the target object and the condition to perform the service (e.g. when the client device is a HP Jornada or a handheld computer with Win CE operation system).



(A)	Content server sends an ISDP request to CSN to subscribe a service for its object. This request is routed to a designated AP server through redirection
(B)	The designated AP server sends an ISDP response back to the content server. The ISDP response contains a service certificate which describes the requested service and rules to perform it as well as a digital signature necessary for performing authentication by the OPES proxy and the AP server
(C)	This certificate is distributed together with the object

(1)	A client sends a HTTP request for that object. The request goes to an OPES caching proxy to see if there is a cache hit
(2)	Assuming no cache hit or the cached version has expired, the OPES proxy forwards the request to the origin content server
(3)	The object along with the certificate is delivered to the OPES proxy
(4)	The OPES proxy sees the certificate and knows the object needs additional services. Based on the information provided in the certificate, the OPES proxy uses the SOAP or ICAP protocol to upload the object to the AP server for value-added services.
(5)	After the AP server performs the subscribed processing, the value-added object with expiration time is sent back to the OPES caching proxy. The OPES caching proxy may decide to cache it
(6)	The value-added object is delivered to the client

**Figure 4.** Post-distribution service performed on behalf of content provider (content server)

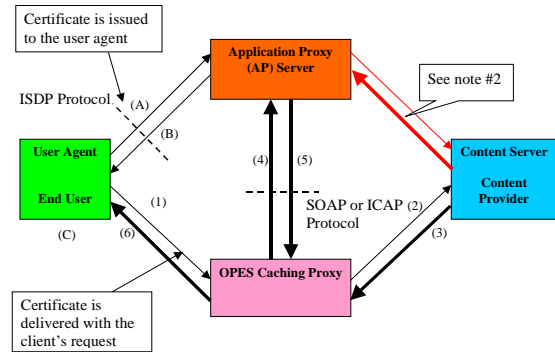
The IETF OPES working group is currently standardizing an XML-based rule specification language for proxy services [3]. This certificate should comply with the format specified by OPES in the future.

In the above example, the OPES proxy requests a video service for the content provider when it receives a client request from a handheld computer. CSN routes the request to one of its AP servers which performs the necessary processing on the video

object and returns the result back to the OPES proxy. The OPES proxy then sends the result to the client, and it may decide to cache the result if appropriate.

### 4.3 Service Performed on Behalf of End User

Figure 5 depicts a typical system interaction for the service performed on behalf of end user. The user agent (or client's proxy) first subscribes to the service from a CSN using the ISDP protocol. Then an application proxy server in the client's proximity responds the ISDP request by sending a certificate back to the user agent (or the client's proxy). This certificate is delivered together with the client's request, instructing other network elements to perform necessary actions on behalf of the end user.



(A)	A client sends an ISDP request to CSN to subscribe a type of service such as adaptive content delivery. This request is routed to a designated AP server through redirection
(B)	The designated AP server sends an ISDP response back to the client. The ISDP response contains a service certificate which describes the requested service and rules to perform it as well as digital signature necessary for performing authentication by the OPES proxy and the AP server
(C)	This certificate is delivered together with the client's request, instructing other network elements the request and response pass through to perform necessary actions on behalf of the client. The certificate may only need to be delivered once to the responsible OPES proxy

(1)	A client sends a HTTP request for that object. The request goes to the OPES caching proxy to see if there is a cache hit
(2)	Assuming no cache hit or the cached version has expired, the OPES proxy forwards the request to the origin content server
(3)	The requested object is delivered to the OPES proxy
(4)	The OPES proxy sees it needs to perform certain actions on that object for the client. Based on the information provided in the certificate, the OPES proxy uses the SOAP or ICAP protocol to upload the object to the AP server for value-added services.
(5)	After the AP server run the subscribed processing, the value-added object with expiration time is sent back to the OPES caching proxy. The OPES caching proxy may decide to cache it
(6)	The value-added object is delivered to the client

- Note:
1. The service subscription (A-C) may be performed by enterprises or ISPs on behalf of their users.
  2. The OPES proxy may skip (2)(3)(4) by asking the AP server to directly fetch the content from the content server. This reduces potential time delay and bandwidth waste for passing big objects back and forth between an OPES proxy and an AP server.

**Figure 5.** Service performed on behalf of end user.

The certificate may only need to be delivered once to the responsible OPES caching proxy. This certificate may contain service instructions at different granularity depending on who initiates the service. If initiated by proxies of enterprises or ISPs, the certificate usually specifies the common services for a group of users. If initiated by end users, it usually specifies the

services for a particular user. An example ISDP certificate issued to a group of end users might look like this:

```
<subscriber = "end user">
  <user ip-address = "15.4.91.*"/>
  <service name = "adaptive and interactive video delivery">
    <rule>
      <property name = "MIME-type" matches =
"video/mpeg">

<action>icap://ms1.csn.hp.com/video_services/summarization_a
nd_keyframe_selection</action>
      <signature>AA%$#&*!@!G%N</signature>
    </property>
  </rule>
</service>
</subscriber>
```

In the above example, a group of end user (e.g. ip-address = 15.4.91.\*) has subscribed to a video summarization service from a CSN. This certificate indicates that video summarization should be performed when a user from the group downloads a MPEG video from the Web.

Although not indicated in Figure 5, the OPES proxy may ask the AP server to directly fetch the content from origin server instead of obtaining the content and passing it to the AP server for value-added services. This reduces the potential time delay and bandwidth consumption for passing a big object back and forth between an OPES proxy and an AP server.

## 5 Redirection

The requests sent to a CSN are redirected to one of the application proxy (AP) servers according to the following consideration:

1. Locality – direct the request to the AP server that is close to it in the network
2. Server load – perform load balancing among the AP servers in a local domain
3. Type of service – direct the request to the AP server that has the corresponding service

Our current redirection scheme is a two-stage process. The first stage deals with network locality (item #1) and is achieved by anycast routing which routes the request to the nearest redirector. The second stage deals with server’s load balancing and service discovery (item #2 and #3) and is performed by the redirector that further redirects the request to an AP server in its control domain. We have studied the problem of redirection in a wide-area network in [24]. We briefly summarize the technique used in our CSN prototype in the following.

### 5.1 Redirection based on Administratively Provisioned Anycast

In a CSN, the redirectors of each domain (see Figure 1) are associated with the same anycast address. Although IP-anycast service is not ubiquitously deployed, it can be provisioned administratively [26]. Assume a CSN’s domain resides in a network domain with A/20 Classless Inter-Domain Routing (CIDR) prefix and an anycast address A\*/32 from the A/20 CIDR block is assigned to the redirectors in the domain. To enable anycast service within the domain, the redirectors

announce host route for shared address A\*/32 through routing daemons such as *gated* [27] or *zebra* [28]. First-hop routers then propagate A\*/32 routes through intra-domain routing protocols such as IGP [29]. With the anycast service provided through IGP, network routers can deliver shared-address A\*/32 packets to the nearest redirector within the domain naturally. The anycast address A\*/32 is advertised to outside domains by BGP [29] together with other A/32 unicast addresses in the A/20 CIDR block. As such, packets sent to A\*/32 from other domains could be folded in with the network route and finally reach the nearest redirector through IGP routing [25].

### 5.2 Redirection based on DNS

After the request reaches the nearest redirector, it is further redirected to an AP server that fulfills the request. We use domain name server (DNS) to perform this second-stage redirection. Essentially, the redirectors are the authoritative name servers for a CSN. This is accomplished by placing DNS “NS” and “A” records in the primary domain name server to identify host with A\*/32 address as the authoritative name server for the CSN’s domain, e.g. “csn.hp.net.” Therefore, when a client requests a service from a CSN’s server, e.g. “ms1.csn.hp.net,” it triggers the client’s local DNS server to issue a recursive DNS query to discover the IP address of “ms1.csn.hp.net.” This DNS query is routed to the nearest authoritative name server (it is also a redirector) through the first-stage redirection. Because informed by the SDM servers, the redirectors have knowledge and information about the AP servers in its domain. After determining which AP server should serve the request, the redirector sends the IP address of the designated AP server as DNS reply to the client. Figure 6 shows the steps involved in our redirection scheme.

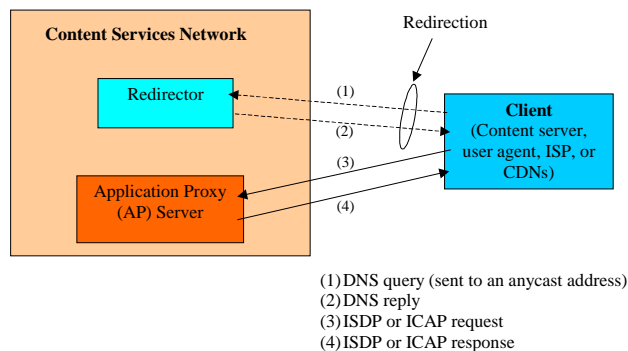


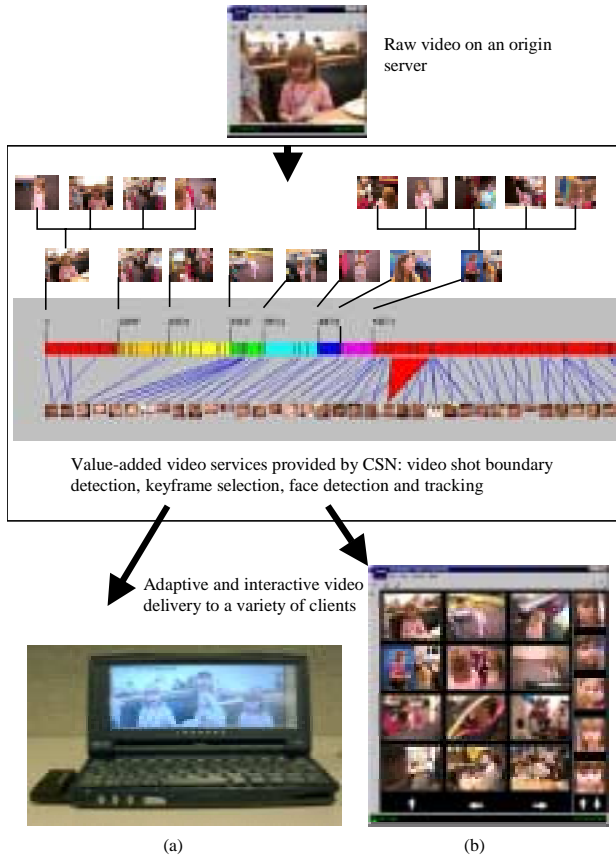
Figure 6. Redirection performed by A CSN’s redirector.

## 6 Current Status

In the process of building a prototype, we decided to use a video analysis technique developed in HP Labs [8][9] as our first CSN’s service. This technique has been used in our previous adaptive content delivery system [13] to perform information abstraction and modality transform. It has also been used in our recent interactive video caching and delivery system [10]. We believe this video service represents a richer set of services that other researchers have not considered [4].

The performed video analysis includes shot boundary detection, keyframe selection and face detection and tracking. The result of analysis is used to enhance the way video is delivered over the Internet, as a value-added service provided by a CSN. For

instance, upon the client's request, the video server can send the meta-data resulted from the analysis that describes the structure and summarization of the video along with the selected keyframe images to the client, so the user can start browsing video immediately without downloading it. In most scenarios, the summary contains sufficient information for the user to decide whether to download or stream any part of a video so that bandwidth and power consumption are saved. This summary also improves the way the user interacts with a video. For example, the user can navigate the content and choose to jump directly to the beginning of a particular shot of interest.



**Figure 7.** As an example service provided by our CSN, we turned a video segmentation and keyframe selection system into an infrastructure service that can be used by content providers or end users to enhance the way video is delivered over the Internet. This example shows the result of video content analysis performed by the application proxy server. The video is segmented into shots that can be managed separately. The keyframes from each shot are selected and clustered to form a hierarchical representation that can be used to provide interactive video delivery. (a) Video delivered to a desktop PC. (b) Video delivered to a handheld computer with wireless LAN connection.

In a bandwidth constraint environment, the above video analysis and the corresponding interactive delivery are valuable services for content providers and end users. Therefore, we have deployed the service in two different modes: mode-I and mode-III as described in Section 4. Based on the service mode-I,

content providers can simply put raw video on their servers and ask a CSN to perform video analysis and host the video for them. The AP servers are equipped with the special video streaming server which can render the video in an adaptive and interactive manner.

Based on the service mode-II, end users (and their corporations and ISPs) can subscribe to the service from a CSN. In this case, the video is first retrieved by the OPES caching proxy and then handed over to a CSN for value-added services. Figure 7 shows the video analysis and the interactive video delivery as a result of the value-added video service provided by our CSN.

## 7 Conclusions and Future Work

Content services network (CSN) broadens the scope of services that have been dealt with by content delivery networks (CDNs) and represents the beginning of an important research area in the future. In this paper, we outlined the architecture of a CSN and the network protocols it uses to interoperate with existing network infrastructures and elements. A new Internet service delivery protocol (ISDP) is proposed to characterize a variety of system interactions necessary to perform different types of service subscription and service rendering. The ISDP protocol supplements many functions not addressed by the iCAP protocol.

We also demonstrated an example service that turned a video segmentation and keyframe selection system into an infrastructure service that can be used by content providers or end users to enhance the way video is delivered over the Internet.

We envision the “plain” content delivery service that CDNs currently offers will soon become a basic service that is no longer a concern. What drives and differentiates the market will be the appearance of new value-added services that fulfill the increasing expectation of Internet users and content providers. Through the development of this new Internet infrastructure, media technologies will find their deserved opportunities to revolutionize the Internet. By making them part of the Internet infrastructure services that can be easily accessed and used like commodity by anyone at anywhere and anytime, content service networks will bring the Internet to the next level that remains to be defined.

Many important and challenging research problems for a CSN remain to be investigated. For example, issues such as security, billing and accounting, service distribution and management are still left open. Our proposed Internet service delivery protocol may need to be revised or extended to encapsulate service scenarios we have not considered. Also, we envision in the future multiple CSNs may coexist and service-peering issues may arise much like the content-peering issues currently happening among CDNs [30]. It is therefore more important that a standardized service delivery protocol is defined and agreed upon. We will continue to work on these issues in the future.

**Acknowledgement:** We are thankful to Shuheng Zhou for the study of redirection problem performed during her internship with HP Labs. We are also thankful for the comments from the reviewers.



## References

- [1] Mark Nottingham, "On defining a role for demand-driven surrogate origin servers," *Computer communication*, vol. 24, no. 2, pp.215-221, Feb. 2001.
- [2] C-H. Chi, Y. Lin, J. Deng, X. Li, T.-S. Chua, "Automatic proxy-based watermarking for WWW," *Computer communication*, vol. 24, no. 2, pp. 144-154, Feb. 2001.
- [3] A. Beck and M. Hofmann, "PSRL: a rule specification language for proxy services," Internet-Draft draft-beck-opes-irml-00.txt, work in progress, Feb. 2001.
- [4] M. Hofmann and A. Beck, "Example services for network edge proxies," Internet-Draft draft-hofmann-esfnep-00.txt, September 2000.
- [5] I. Cooper, I. Melve, and G. Tomlinson, "Internet web replication and caching taxonomy," Internet-Draft, RFC-3040, January 2001.
- [6] G. Tomlinson, H. Orman, M. Condry, J. Kempf, and D. Farber, "Extensible proxy services framework," Internet-Draft draft-tomlinson-epsfw-00.txt, July 2000.
- [7] J. Elson et al. "ICAP the Internet content adaptation protocol," Internet-Draft draft-elson-opes-icap-01.txt, Aug 2001.
- [8] W. Y. Ma and HongJiang Zhang, "An indexing and browsing system for home video," X European Signal Processing conference (invited paper), Finland, September 2000.
- [9] HongJiang Zhang and W. Y. Ma, "Structured and content-based video browsing," Proceedings of the 32th Asilomar Conference on Signal, System & Computers, Pacific Grove, CA, 1998.
- [10] S. J. Lee, W. Y. Ma, and B. Shen, "Interactive video caching and delivery using video abstraction and summarization," Proc. International Workshop on Web Caching and Content Distribution (WCW'01), June 2001.
- [11] Randy H. Katz, "The Post-PC Era: It's All About the New Services-Enabled Internet," a presentation given to HP Labs in 2000, <http://www.cs.berkeley.edu/~randy>.
- [12] A. Fox, S. D. Gribble, Y. Chawathe, and E. A. Brewer, "Adapting to network and client variation using active proxies: lessons and perspectives," *IEEE Personal Communication*, Vol. 5, No. 4, pp. 10-19, August 1998.
- [13] W. Y. Ma, I. Bedner, G. Chang, A. Kuchinsky, and H.J. Zhang, "A framework for adaptive content delivery in heterogeneous network environments," *SPIE Multimedia Computing and Networking 2000*, pp. 86-100, San Jose, January 2000.
- [14] Spyglass-Prism. <http://www.spyglass.com>.
- [15] Intel QuickWeb. <http://www.intel.com/quickweb>.
- [16] ProxiNet. <http://www.proxinet.com>.
- [17] O. Angin, A.T. Campbell, M. E. Kounavis, and R. R.-F. Liao, "The Mobeware Toolkit: Programmable support for adaptive mobile networking," *IEEE Personal Communications*, Vol. 5, No. 4, August 1998, pp. 32-43.
- [18] T. Bickmore and B. Schilit, "Digestor: Device Independent Access to the World Wide Web", proceedings of the Sixth International World Wide Web Conference, Santa Clara, California, 1999.
- [19] B. D. Noble, M. Satyanarayanan, D. Narayanan, J. E. Tilton, J. Flinn, and K. R. Walker, "Agile application-aware adaptation for mobility," Proceedings of the 16<sup>th</sup> ACM Symposium on Operating System Principles.
- [20] J. Smith, R. Mohan, and C. Li, "Scalable multimedia delivery for pervasive computing," *ACM Multimedia*, 1999.
- [21] C. Yoshikawa et al, "Using Smart Clients to Build Scalable Services," Proc. Winter 1997 USENIX Tech. Conf., January 1997.
- [22] The Internet Content Adaptation Protocol (iCAP). <http://www.i-cap.org>.
- [23] Open Proxy Extension Services. <http://www.cs.utah.edu/~horman/opencache.html>.
- [24] Shuheng Zhou, "Wide-area redirection," HP Technical Report, 2000.
- [25] C. Partridge, T. Mendez, and W. Milliken, "Host anycast service," RFC-1546, November 1997.
- [26] Steve McCanne, "Enabling Internet TV With Intelligent Network Infrastructure," mcast 2000, San Francisco, February 2000.
- [27] Merit GateD Consortium: <http://www.gated.org>.
- [28] GNU Zebra – routing software: <http://www.zebra.com>.
- [29] Christian Huitema, *Routing in the Internet*, Prentice Hall PTR, 2000.
- [30] Content Bridge Alliance. <http://www.content-bridge.com>.
- [31] Simple Object Access Protocol (SOAP). <http://www.w3.org/TR/SOAP>.