Towards a Peer-to-Peer Extended Content Delivery Network

Daniel Pakkala and Juhani Latvakoski
VTT Electronics, Kaitoväylä 1, P.O. Box 1100, FIN-90571 Oulu, Finland
E-mail: {Daniel.Pakkala, Juhani.Latvakoski}@vtt.fi

Abstract—**The contribution of this paper is a peer-to-peer (P2P) extension concept for star topology content delivery networks (CDN). The P2P extension concept targets enhancement of the functionality and extends the reach of the legacy star topology-type of CDNs by means of P2P techniques. Various P2P technologies are reviewed in the context of content delivery platforms and reasoning for the concept is provided. Finally, a set of application scenarios for the utilization of the P2P extension concept is provided.**

Index Terms—**Content Delivery Networks (CDNs), Peer-to-Peer CDN Extensions.**

I. Introduction

A well-known concept of content delivery networks (CDN) is based on the idea that no content delivery infrastructure is needed from the content providers, but the content distribution and delivery is outsourced to CDN providers that have business relationships with the content providers [1]. There are many different types of CDNs, depending on the network topology and business relationship. In the star topology CDNs, like Akamai [2], the content is inserted into the center of the network, and further distributed and replicated into caches of specialized edge servers residing close to the end users. The content is then delivered to the end users from the edge servers. Caching or broadcast technologies can be used for content distribution. The edge servers are typically located inside the networks of the access connection provider, such as an Internet Service Provider (ISP), and can act as a caching proxy or receive the content they host from, e.g., satellite broadcast. Such a legacy star topology CDN is illustrated in Figure 1a.

In addition to replicating the content to the edge of the network, the CDNs incorporate a request-routing system that routes the client requests to the optimal edge server replica. The request routing system can utilize information about geographical location, network conditions and load on the edge servers to select the best possible edge server replica for a given client. The value of a CDN can be determined by its scale and reach [1] - the scale refers to size of the aggregate infrastructure and the reach refers to the increased diversity of the content locations. The content replicas on the edge servers close to the end users overcome the issues of network size, network congestion and network failures in CDNs. Another advantage of using CDNs is that no additional and expensive infrastructure is needed from the content providers as the content distribution and delivery can be outsourced to CDN providers. Another, essentially different, type of CDN is emerging: the P2P CDN.

Acknowledgement: This work has been partially funded by the European Union Information Society Technology in the Sixth Framework Programme under contract number FP6-507993 (mCDN). As such, the paper may not reflect the strategies of companies involved in the mCDN project. We thank our project partners for their good co-operation in the project.

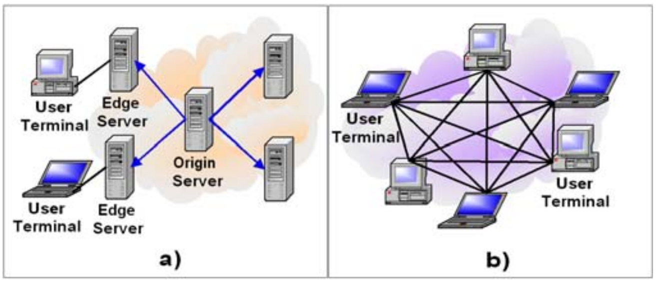

Fig. 1. Two different CDN architectures illustrated. On the left side, (a) a star topology CDN architecture, and on the right side (b) a P2P CDN architecture.

Usually, the term Peer-to-Peer (P2P) refers to a specific type of distributed system, which does not rely on traditional client/server architecture but uses the servents-based architecture. A servent includes both the client and the server. In addition, P2P systems usually do not depend on any centralized point of control, even if they are not excluded. Some P2P solutions apply special servers, also referred to as super peers, as centralized control points for such functions as P2P network formation, content indexing and searching. This divides the P2P architecture solutions into two different groups; hybrid and pure P2P architectures [3]. Sometimes, a subclass of hybrid P2P architectures, referred to as server-centric P2P architecture, is described as its own type of architecture.

In the P2P CDNs the content is delivered straight from and to the end user terminals that constitute the delivery network. The advantages of P2P CDNs include independency of central servers enhancing system robustness and network topology enabling more balanced network utilization than in the client-server architectures. Figure 1b. illustrates a pure P2P CDN consisting of user terminals.

The P2P technology-based content delivery systems have obvious advantages regarding load balancing, dynamic information repositories, fault tolerance, availability, content-based addressing and improved searches [4]. However, the pure P2P model is not an applicable architecture choice for a commercial CDN as such. As the contribution of this paper, we propose a P2P Extension concept to the star topology CDN architecture to gain most of the advantages provided by P2P technology without the known drawbacks.

The rest of this paper is organized as follows. Chapter 2 reviews P2P technologies in the context of content delivery platforms and provides the reasoning for the P2P Extension concept, which is described in Chapter 3. An overview set of application scenarios for the P2P Extension concept is provided in Chapter 4. Finally, some concluding remarks are provided in Chapter 5.

II. Review of Peer-to-Peer Technologies

The P2P system architectures can be roughly divided into hybrid and pure P2P architectures [5], [3]. Sometimes, a subclass of the hybrid P2P architectures, referred as server-centric P2P architecture, is presented as its own type of architecture.

Hybrid P2P architectures utilize specialized servers for some operations, like indexing and resource and peer discovery, and are dependent on the specialized servers. In hybrid P2P architectures the peers may first interact with the specialized servers to gain knowledge of other peers and resources in order to directly share data or resources with them. Examples of this kind of server-centric architecture are Napster [6] SETI@home [7], and Groove [8].

The other P2P system architecture type, commonly referred to as the pure P2P model, does not incorporate any specialized servers. Instead, all the peers are treated as equal in the P2P system. Gnutella [9] is an example of a pure P2P system. In a pure P2P architecture the peers use distributed search mechanisms to search for other peers and resources inside the P2P system and share data or resources directly with other peers.

The P2P system architectures are often based on overlay protocols that form communicating layers across the member peers of the P2P network. The overlay protocols typically rely on existing communication protocols like TCP, UDP and HTTP over the IP. Often, the lowest overlay protocols of the P2P systems, like the ConstructOverlay protocol in [10] and the Peer Discovery Protocol of JXTA [11], are used for finding or sharing information about the other peers, peer groups or network structure in general when a new peer joins the P2P network. Furthermore, the P2P architectures often incorporate separate protocols for additional communication overlays, indexing, search and delivery of data.

The legacy client-server model is used for implementing many shared applications with a successful outcome [5]. However, these applications are highly dependent on centralized servers, reducing the scalability and reliability of the applications. In client-server applications the load of the server increases linearly in proportion to the number of clients and the system is vulnerable to server failures, leading to failure of the whole system. The hybrid P2P architectures have many of the same system reliability weaknesses as the client-server architectures as they are often dependent on specialized servers for peer and resource discovery. The pure P2P architecture is free of the aforementioned weaknesses because it does not incorporate specialized servers. However, the cost is that more complex distributed search mechanisms are needed, which generates more network traffic. And the mechanisms needed for establishing the network in pure P2P architectures are much more complex than in the hybrid P2P architectures.

The advantages of the P2P-based architectures include enhanced load balancing, dynamic information repositories, high fault tolerance and availability via redundancy, content–based addressing and improved searches [4].

A. *Reasoning for the Concept*

The pure P2P CDN approach has many issues that make the position of the P2P CDN providers unstable. First, the fact that the pure P2P CDN infrastructure is based on dynamically joining and leaving end user terminals with relatively slow network access connections reduces its performance when compared with legacy star topology CDNs. Second, the fact that users volunteer to provide their terminals and access connections as resources for the P2P CDN makes the value parameters of a CDN, the scale and reach, fully dynamic. Third, the heavy bi-directional bandwidth load that the P2P systems generate for the ISPs, carriers and other connection providers [12], [3] indicates that in the long run the CDN infrastructure of the P2P CDN providers may be in danger as the volunteer infrastructure lending, with no return value, is no longer considered reasonable by the connection providers. The increased user traffic may also cause network congestion, forcing the connection providers to consider capital investments to speed up their backbone networks [3]. Finally, the fact that the content is stored in end-user terminals poses very strict requirements for the encryption and DRM issues in P2P CDNs, if used commercially to compete with legacy CDNs, when unauthorized access to content must be prevented. Due to all the aforementioned issues, the emerging pure P2P CDN model is not expected to have very much potential for competing with legacy star topology CDNs in a commercial CDN context.

As the drawbacks of star topology CDNs, like Akamai [2], are the huge maintenance and infrastructure costs [13], and the drawbacks of pure P2P CDNs include the instability of the CDN infrastructure and dynamic volunteer resources [13], a compromise hybrid solution between a static server infrastructure and dynamic volunteer P2P resources provides an interesting approach to constitute a P2P enhanced and extended hybrid CDN. In this paper this approach is referred to as the concept of P2P Extended CDN. In this hybrid solution the static CDN infrastructure is used to distribute content from the originating server to static edge servers that further replicate and route requests in the dynamic P2P networks on the network edge. This approach targets the reduction of the infrastructure costs of the CDN provider and facilitates the handling of flash crowds on edge servers [10] via enhanced load balancing and intelligent replication techniques. If the P2P extension nodes are also equipped with applicable caching mechanisms that lower the bandwidth consumption, the P2P Extended CDN approach should also be acceptable for the connection providers. Providing some kind of return value to volunteer P2P participants might also enhance the approaches available to a dynamic P2P infrastructure at the network edge. This return value to volunteer P2P resource providers, like P2P communities or CDN service providers, could be to support various P2P services in addition to content delivery at the CDN edge. Supporting various content-based P2P services is part of the P2P Extended CDN approach.

III. P2P Extension Concept for CDNs

The architecture of the envisioned P2P Extension concept is visualized in Figure 2. In the concept the star topology CDN, consisting of distribution and edge servers, is extended with P2P overlay networks at three different domains. Each domain is considered its own P2P group in the P2P Extension concept.

First, the infrastructure of the CDN service provider is extended with P2P-connected edge servers, introducing enhancements regarding load balancing, fault tolerance and availability via redundancy, content-based addressing, and improved searches. The second P2P extension is a P2P-connected available-computing infrastructure in the CDN service provider or customer domain providing additional dynamic information repositories for content replicas and further enhancing load balancing and content availability. Finally, the third possible P2P extension is the P2P-connected end user terminals at the end user domain, enabling provisioning and use of various P2P multimedia services. As illustrated in Figure 2, the architecture of the P2P extensions is hierarchical in the sense that the P2P-connected edge servers host and manage the extensions in the CDN service provider domain that further host and manage the possible end user domain extensions.

The central focus of the concept of the P2P Extended CDN is how to enhance the star topology CDN networks with managed and dynamic P2P extensions that would reduce the need for a static edge server infrastructure, enable efficient load balancing and boost up the delivery of large files on the "last mile" of delivery without the known problems related to the pure P2P-based content delivery systems.

The concept of the P2P Extended CDN considers the utilization of P2P technologies in a managed and secure manner, providing no user anonymity. Accordingly, the peers wishing to become members of the dynamic P2P extensions at the CDN edge are authenticated and authorized, and therefore considered to be members of a trusted domain that can be managed by the P2P or CDN service provider.

In the P2P technology-based CDN at least some content replicas are usually stored in the member peers. This gives some advantages but also creates some security concerns as these peers might be the end user terminals. However, in P2P technology-based content delivery systems other entities can also act as peers of the P2P system. For example, in the Joltid PeerCache™ [14] solution the peers are caching proxies that communicate with each other in a P2P manner to avoid centralized points of network load on the servers. If the content is stored outside the trusted domain boundaries by any P2P-based CDN, the content must be made unusable to unauthorized usage by using various content encryption technologies. Figure 3 illustrates the domain boundaries of the envisioned P2P Extended CDN approach.

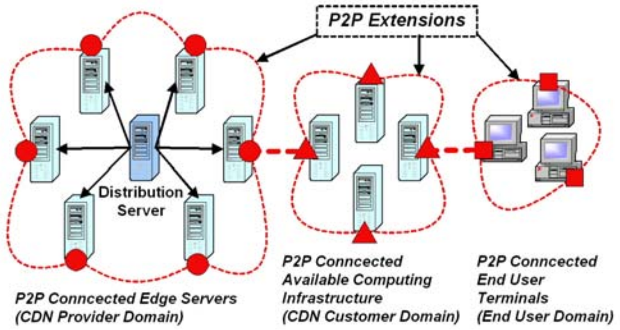

Fig. 2. P2P extended CDN architecture. The figure illustrates the three different hierarchical levels of the envisioned P2P extensions of a star topology CDN.

The static star topology CDN is illustrated on the right side of Figure 3. On the left side the dynamic P2P extension network with peers located at a trusted domain - such as a connection to a provider network (triangles) - and at a distrusted domain - such as unknown subnets (squares) - is illustrated. These together form the envisioned P2P Extended CDN addressing end-to-end content distribution and delivery from the origin server to the end user terminal.

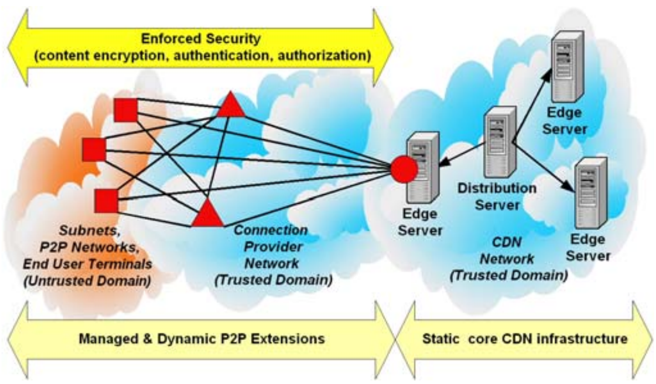

Fig. 3. Domains of the P2P extensions. The figure illustrates the different network domains where the P2P extensions are envisioned to be deployed.

The concept of the P2P Extended CDN also targets the optimization of the “last mile” of content delivery from the edge server to the end user terminal, which is mostly left unattended by the legacy CDNs. The “last mile” delivery of content could be optimized using parallel downloads from neighboring peers, reducing download times and alleviating the load on the single edge servers hosting popular content. This further enhances and complements the performance of the star topology CDNs that are usually targeted at overcoming network-related issues like network size, network congestion, network failures and load balancing between the origin server and the edge, leaving the “last mile” of content delivery from the edge server to the end user terminal with little attention. With the P2P enhancements that further replicate content and enable parallel downloads at the network edge, the “last mile” of content delivery can be optimized and overall content access times can be reduced. The approach could “boost up” the delivery of very large files like video or audio files. This would definitely add to the value of the CDN as it would contribute to the value parameters [1] of scale and reach. The scale parameter, referring to the aggregation infrastructure of the CDN, as well as the reach parameter would be enhanced by the dynamic P2P extensions by further replicating the content on the edge of the network, thus enhancing content availability. Further advantages of P2P extensions may include more balanced network loads at the connection provider’s domain.

Having a static core CDN infrastructure into which the dynamic P2P extension network is “hooked” alleviates the problem of high bi-directional bandwidth requirements into and out of the connection provider’s domain, which is a problem identified in almost all of the existing P2P-based content delivery systems [12]. In the P2P Extension concept this problem is avoided as the core CDN infrastructure distributes the content from the origin servers to near the user, which usually means to the connection provider’s network, instead of downloading the content directly from peers belonging to some other network domain, which is the case in many P2P-based content delivery systems today. In the envisioned P2P Extended CDNs the edge servers could efficiently cache the content if direct content delivery between the peers belonging to different subnets is blocked or delivered along a certain known caching policy.

By introducing the “always present” edge servers of the star topology CDN as super peers managing the P2P extension clouds, stability and a central point for authentication and authorization can be introduced to the P2P Extended CDN. This enables the dynamic P2P-based part of the delivery network to be managed. The setup is similar to restricted membership P2P communities, where security can be enforced. As super peers, the edge servers of the static core CDN could also monitor the usage of the content centrally, and thereby provide the means for charging the users for the commercial content consumed.

IV. Application Scenarios

The application scenarios of the P2P Extended CDN concept include all the content distribution and delivery scenarios of legacy star topology CDNs with P2P-enhanced load balancing, information repositories, fault tolerance and availability. In addition to these application scenarios the P2P Extended CDN concept can be applied to enable end user-created content provisioning, group and community content distribution and delivery, and various P2P services.

The amount of digital content created by users or consumers is growing all the time as digital cameras, video cameras and other recording devices become more commonplace and are integrated into mobile devices. A legacy star topology CDN addresses the content delivery problem from the network core to the network edge, but not other way around - as is the case regarding end user content delivery. The concept of a P2P Extended CDN could enable the end users to act as content providers by supporting P2P content services that enable delivery of content to other end users and user groups. Here the content delivery would take place at the network edge, utilizing any commercial P2P Extension infrastructure.

The P2P services supported by the concept of a P2P Extended CDN can be virtually any services incorporating multimedia communications between peers. Additional P2P services that are envisioned for the P2P Extended CDN concept are, e.g., supporting community multimedia content delivery services, various content messaging services, multimedia content storing and management, and P2P workspace hosting.

Part of the services provided for the “P2P infrastructure lenders” could be free of charge to boost up the available P2P resources; additional P2P services might be chargeable. Accordingly, the P2P services considered in the concept of a P2P Extended CDN could bring new business opportunities and revenue for CDN service providers if the services are commercially provided with an appropriate business model consisting of actors like content provider, CDN provider, CDN service provider, P2P service provider, and the end user, either in a consumer or provider role.

V. Concluding Remarks

In this paper we have presented the concept of a P2P Extended CDN. A review of P2P technologies in the context of content delivery platforms and the reasoning for the P2P Extension concept has also been provided. The concept itself has been described and various application scenarios for the P2P Extension concept briefly overviewed. Topics requiring future research include P2P service provisioning, design of a modularized architecture with well-defined inter- and intra-layer interfaces, and validation of the concept via prototype implementations.

REFERENCES

[1] M. Day, B. Cain, G. Tomlinson, and P. Rzewski, "A Model for Content Internetworking (CDI)," IETF, RFC 3466, URL: http://www.faqs.org/rfc/rfc3466.txt, February 2003.

[2] Akamai, "Web Site," URL: http://www.akamai.com, 2005.

[3] T. Oh-ishi, K. Sakai, K. Kikuma, and A. Kurokawa, "Study of the Relationship between Peer-to-Peer Systems and IP Multicasting," in *IEEE Communications Magazine*, vol. 41, 2003, pp. 80-84.

[4] M. Parameswaran, A. Susarla, and A. B. Whinston, "P2P Networking: An Information-Sharing Alternative," in *IEEE Computer*, vol. 34, 2001, pp. 31-38.

[5] M. Nakamura, J. Ma, K. Chiba, M. Shizuka, and Y. Miyoshi, "Design and implementation of a P2P shared Web browser using JXTA," presented at Advanced Information Networking and Applications, 2003. AINA 2003. 17th International Conference on, 2003.

[6] L. Napster, "Web Site," URL: http://www.napster.com, 2005.

[7] SETI@home, "Web Site," URL: http://setiathome.ssl.berkeley.edu/, 2005.

[8] GrooveNetworks, "Web Site:," URL: http://www.groove.net/, 2005.

[9] Gnutella, "Web Site," URL: http://www.gnutella.com/, 2005.

[10] A. Stavrou, D. Rubenstein, and S. Sahu, "A Lightweight, Robust P2P System to Handle Flash Crowds," *IEEE Journal on Selected Areas in Communications*, vol. 22, pp. 6-17, 2004.

[11] R. Flenner, M. Abbott, T. Boubez, F. Cohen, N. Krishnan, A. Moffet, R. Ramamurti, B. Siddiqui, and F. Sommers, *Java P2P Unleashed: With JXTA, Web Services, XML, Jini, JavaSpaces, and J2EE*, 1 ed: Sams Publishing, 2002.

[12] S. Saroiu, K. P. Gummadi, R. J. Dunn, S. D. Gribble, and H. M. Levy, "An Analysis of Internet Content Delivery Systems," presented at OSDI '02: 5th USENIX symposium on Operating systems design and implementation, Boston, MA,USA, 2002.

[13] R. Wetzel, "CDN Business Models -The Drama Continues," *Business Communications Review*, pp. 51-57, 2002.

[14] Joltid, "PeerCacheTM Web Site," Joltid Ltd. URL: http://www.joltid.com/index.php/peercache/, 2005.