

Internet Content Delivery Systems– Growing the Internet Intelligently

Within a few short years, the Internet has become an essential part of our daily lives, providing email, shopping, research, and news. Approximately 37 million U.S. adults use the Internet at home on a daily basis, compared to only 19 million in mid-1997. The number of daily Internet users at work has also skyrocketed from 19 million to 32 million since mid-1997.¹

The Growing Internet Masses

Intense competition for Internet access has driven the costs down dramatically to the point where consumers can get unlimited Internet access for as little as \$20 per month. Internet Service Providers (ISPs) are struggling to attract as many users as possible by providing better service, keep costs down and profits up.

Keeping the quality of service up is a difficult task due to over-congested networks that provide agonizingly slow access especially during peak times. Although delivery of email messages and Web content is adequate, the quality of service for those with Internet access speeds greater than 28.8 Kbps is lacking. Today's modems are promising end users access speeds greater than 56k; yet in reality, current technology in many areas of the country does not support this speed.

Internet Service Providers are rapidly introducing DSL and cable modem technology, offering premium priced, high-speed (128Kbps to 1.5Mbps) bandwidth to the home and small businesses. Although these technologies promise faster access and richer services such as video on demand, DSL users can easily overwhelm the provider's upstream network connection unless additional expensive bandwidth is purchased. Even then, "flash" crowds can suddenly request the same content simultaneously, causing a complete shutdown in access to these sites. Some kind of buffering is essential to deal with these events.

Content providers of today's web sites are limited not by their imaginations, but instead by the inability of networks to handle richer content. Designers simply can't use streaming digital video, digital audio, large and flashy graphics and other tools that will easily overwhelm a typical user's computing access. A content provider will generally restrict their Web page content to achieve optimal delivery speeds of 29.8 or less -- primarily because of network bandwidth costs and the unsuitability of today's Internet for delivering higher bandwidth content.

Increasing Bandwidth Availability

The industry is responding to this challenge with two fundamental strategies: building bigger and faster networks and delivering content in more intelligent ways.

Although larger networks -- complete with terabit routers and wave-division multi-plexing (WDM) -- are currently being built between heavily congested networks, this expensive alternative does not completely solve the problem. The content still has to be routed over these networks, sometimes manually, through peering arrangements or other schemes. As a result, today there are brand new, high-speed networks sitting around totally quiescent, because traffic has no way of being directed over them.

A better solution to the bandwidth bottleneck involves delivering content in more intelligent ways. Internet providers are now improving their end-users' experience by using new 'caching' techniques. A cache is a server that replicates, stores and delivers Internet content, such as Web pages, files and documents, or streaming media such as Real Audio or NetShow. Caches are strategically placed close to the user or requester of the content (and/or at strategic aggregation points) in order to eliminate potential bottlenecks. Internet users are getting content from caches on local networks -- which are faster, more reliable, and

¹ Internet User Trends: Year-End 1998, a biannual study produced by The Strategis Group.

consume less bandwidth than accessing a remote origin server. Better still, most caches are deployed transparently, so that neither the content provider nor the end-user has to do much configuration to use content served from cache.

(graphic from web basics white paper)

Today's caching techniques speed up the user's Internet experience in poorly connected rural and remote areas where upstream bandwidth is scarce and expensive. In better-connected networks, caches provide outstanding buffers to handle "flash crowds" such as Victoria's Secret broadcasts and on-line Clinton testimony. Data served from cache avoids the congestion caused by massive duplicate requests from users for the same content.

Caching sounds like an ideal solution to the bandwidth dilemma, but there is a downside. Because today's first-generation "dumb" caches are transparent to both the user and the content provider, the cache has to guess or invent schemes to maintain integrity of the content. For example, a cache needs to check for possible discrepancies between its copy of the content and the original. This is called "cache coherency".

A recent university study determined that, in well-connected networks, cache-coherency checks can consume 40%-60% of available bandwidth, offering a net reduction of overall bandwidth!²

Because of these problems, today's caching solution is really not a solution at all.

A Better Solution – Intelligent Content Delivery

Internet Content Delivery Systems (ICDS) offer a better, more extensible way for service providers, content providers, and service consumers to grow their Internet businesses. ICDS products move Internet content closer to users through "intelligent" hardware and software products, thereby substantially reducing Internet traffic congestion and enabling new applications.

Entera, Inc., located in Fremont, California, develops ICDS for Internet network operators and content providers. The benefits of ICDS include substantially reduced costs, an enhanced user experience, and overall reduction of Internet congestion for companies.

The "engine" for building an ICDS is TeraNode, a standards-based, scalable software platform. TeraNode is the world's first intelligent, multi-protocol, Internet Content Delivery System. It works with any Internet network and can be deployed at remote Points of Presence (POPs) or central data centers. TeraNode is compatible with a wide range of Internet protocols, and provides an innovative way to cache and deliver frequently used content--including Web, audio/video and USENET news. TeraNode is extremely cost-effective for small POPs and can handle high-end, clustered applications. TeraNode's open API allows additional protocol and content types to be supported and provides for easy third party application development.

How TeraNode Works

TeraNode allows ISPs and corporations to speed up Internet response times by moving critical content to where it is needed -- at the edge of the network close to its users. The TeraNode software/hardware system sits at a POP close to the end users. TeraNode caches and replicates content that is either requested by users (using transparent caching) or pushed by content providers. The result is that users get faster access to frequently used content without consuming upstream bandwidth.

² Web Cache Coherence, Adam Dingle, KSVI Charles University, Prague, Czech Republic, Tomas Partl, FJFI Czech Technical University, Prague, Czech Republic

Intelligent caching is a significant step forward in managing Internet traffic. The new TeraNode ICDS caches and serves HTTP, streaming media, and USENET news faster and more effectively than ordinary cache or mirroring products.

The two keys to TeraNode's intelligent caching capability are its *open API* and its *relational database engine*. TeraNode's Open API lets ISPs, content providers, and third-party developers build custom applications to replicate, push, expire, charge for, secure, screen, and otherwise manage Internet content intelligently. The Relational Database Engine stores every object--every web page, every graphic, every movie, and every audio stream--as a database record, along with critical information about the object.

These unique features let TeraNode make intelligent caching decisions about Internet content and give ISPs and content providers the flexibility to define policies, add new protocols, develop new applications, and implement business systems.

(graphic of product from product sheet)


TeraNode is a rack-mountable hardware and software solution built specifically for commercial applications and is expandable not only via software plug-ins but also via hardware expansion and upgrades. As a result, TeraNode can be economically deployed at any point in the network similar to a router or switch. This is a more cost-effective solution compared to expensive, high-end centralized caches and servers. TeraNode's open design enables companies to leverage existing infrastructure investments, quickly and inexpensively adding the capability they need -- when they need it.

Real World Applications

TeraNode makes it easier to deliver high-speed access and value-add services that can increase business while reducing bandwidth and administration costs.

Adding DSL Services

ISPs are currently adding DSL services, cable, or other high-speed access that place additional strain on existing backbone and upstream connections. Instead of spending time, money, and staff redesigning the backbone to increase bandwidth, TeraNode offers a better solution.



(photo from web page -- solutions--DSL)

Usage patterns change significantly when high-speed services are added. With 192Kbps or 1Mbps bandwidth, high-speed end users can put an unpredictable load on the network. By placing a TeraNode in POPs close to DSL users, ISPs can offer more connections without major increases in upstream bandwidth. TeraNode also acts as a buffer for massive downloads and "flash" traffic such as widely-publicized and promoted Internet events.

A router can be set to redirect HTTP traffic to the TeraNode, making it transparent to the end users. As users request Internet objects (Web pages, streaming video, news, etc.), TeraNode checks its cache. If the objects are in cache, TeraNode serves them directly to the users without using upstream bandwidth. If the objects are not in cache, TeraNode fetches them from their source and stores them in cache for later retrieval.

Popular sites and major download "burst" events are handled gracefully in this manner. TeraNode is the only cache architected to handle not only HTTP, but also FTP, streaming media, and other bandwidth-hogging datatypes. More importantly, as DSL service becomes more widely deployed and competitive, ISPs can use TeraNode as an extendable platform to deliver value-add services using this new added bandwidth.

Eliminating News Headaches with TeraNode


Many ISPs provide news service to their subscribers. A complete set of news feeds can easily drive between 14-20 gigabytes a day of traffic through an upstream connection, which can completely saturate a T-1 or E-1 uplink. At that point, the ISP needs to find room to store 90 days of articles (which could add up to a terabyte or more) or offer restricted news groups or more-frequent expiration. These news-feeding and news-storing hassles have caused many ISPs to outsource their news service, which, although quite convenient, adds cost and can still consume a fair amount of upstream bandwidth.

A better way to deliver news service is to let TeraNode cache and serve USENET news. TeraNode satisfies even the most demanding news users quickly, reliably, and with much less maintenance than ordinary news server software. About 80% of news traffic comes from the top 100 groups, which can be pre-populated or "pushed" into the TeraNode using its unique LiveFeed feature, giving users quicker response times. Also, news articles are broadcast once and don't change afterwards, which makes them inherently cacheable. TeraNode subscribers simply point their newsreaders to the TeraNNRP software running on the TeraNode server. TeraNNRP pulls the requested overviews and articles off of the main news and stores the data locally so that the next user requesting the same article is served directly from the TeraNode. This process saves upstream bandwidth and the connection charges from the news service.

Using Satellite Feeds to Push Content

Caches are normally populated with content when users access them. The cache fetches the objects over an upstream link and stores it. This is sometimes referred to as "pulling" content through the cache. Assuming the content doesn't change, it is still available in the cache for the next user who requests it.

As more users "pull" content through the cache, the content is stocked more quickly and response times improve. But if a small subscriber base isn't pulling much traffic through a cache, it does not get filled as quickly with popular content. TeraNode solves this problem by tightly integrating with satellite services such as SkyCache to proactively push popular content into the cache. This capability results in faster to the ISP or organization.



The Future of High Speed Internet Access

Internet usage is growing at an extraordinary pace as more and more people daily utilize this indispensable resource. High quality graphic, audio and video delivery over the Internet is just around the corner. Internet Telephony, Video Conferencing and other alternative Internet applications are on the sidelines waiting for mass access to high speed networking technology. This exciting Internet explosion demands new and innovation methods of dealing with the bandwidth challenge.

The most effective way to deal with this challenge is through Intelligent Content Delivery Systems. Entera Inc.'s TeraNode is the world's first intelligent, multi-protocol, Internet caching node. This "intelligent" hardware/software product offers substantially reduced costs for ISPs, faster Internet access for the end-user, and an overall reduction of Internet congestion for businesses.

[Sidebar]

TeraNode – Key Features

- Add high-speed services such as DSL while minimizing upstream costs
- Install inexpensively and expand easily as your network grows
- Develop and add new services easily with third-party or your own plug-ins
- Integrate with your business and user management systems via RADIUS or SQL
- Interoperates with your existing routers, caches, servers, and services
- Avoid maintenance hassles with easy installation, proven and reliable hardware, fail-safe design and commercial support.

[word count: 2,152]