

In Search of the Intelligent, Profitable, Network Edge

Net Forecasts – Peter J. Sevcik

BCR Volume 32, Number 5

May 2002

A year ago, I described an idea for a new application platform that I called the Network Palette (see *BCR*, May 2001, pp. 10-11). I argued that highly distributed content-delivery networks would morph into highly-distributed computing and storage platforms, from which true network-based applications would be delivered to users. The Network Palette was based on a model that was midway between the centralized datacenter model and the completely decentralized desktop model.

Well, the time has come to ‘fess up: My model for an intelligent network edge is not going to flourish. However, other intelligent edge plays will succeed.

The Datacenter Edge (Centralized)

One “edge” where big things are happening is in datacenters, where computing is segmenting into three components: computing, storage and networking. The first component to separate itself from an integrated computing model was storage, and today an entire industry segment has emerged, producing devices that off-load storage, file management and redundancy tasks from the application and database computers.

This trend is also extending into networking--onerous and interrupt-driven network functions are being ported to separate, special-purpose boxes and appliances (see this column in January 2002, pp. 12-13). This may be a temporary phenomenon, however, as many functions will probably re-integrate themselves into next-generation front-end processors (FEPs).

It’s clear why reintegration is likely: There are simply too many of these devices now, as shown in the list below:

- Caching
- Class-based queuing
- Denial of service protection
- DNS proxy
- Firewall filtering
- Flow measurement and reporting
- Global load balancing
- Intrusion detection

- IPSec encryption
- Key management
- Local load balancing
- Rate shaping
- Route control
- SSL acceleration
- TCP flow acceleration
- TCP termination and multiplexing
- Traffic grooming
- User authentication
- XML pre-processing
- Policy control for all of the above

This next-generation FEP will handle complex processing between the outside world (WAN) and the datacenter, and operate at many protocol layers. The intelligent edge in this vision is *network computing* at the edge of the datacenter; computing for the benefit of the network.

Once network computing catches on, next-generation IP-based FEPs will probably be given a new, catchier name--like IP Service Machines (ISMs). ISMs will continue to be required for service delivery by both public and enterprise datacenters, and some scaled-down versions will migrate directly to enterprise locations.

The Wide-Area Network Edge (Modest Distribution)

The wide-area network (corporate or public Internet) is a transport system. It can perform some of the datacenter, network-computing functions listed above, but it also has unique responsibilities that are within the domain of the service providers--access, switching, billing and traffic-engineering functions that are needed to move packets across networks.

However, it’s easy to envision traditional servers or even the pure storage or computing components being put into many locations on a transport network to perform non-transport functions. The first implementation of this idea occurred on content delivery networks (CDNs), which use servers dispersed across the Internet to provide their unique functions. This is the *computing*

network vision of the world; multi-location computing for the benefit of the application.

One of the key issues in CDNs is scalability; operators have to decide whether to deliver their service out of fewer, but very large locations, or to deploy a very large number of serving locations. Mirror Image, Equinix and AT&T have taken the former approach--it's essentially a datacenter model for CDNs. They provide CDN functionality out of about 15 datacenters, strategically located around the world. All three claim to be only one autonomous system (AS) hop away from 90 percent of the world's Internet population.

Akamai is the only CDN that takes the latter approach--it operates 13,500 servers deployed in more than 1,000 locations. Akamai can probably make a much better claim about being closer to even more than 90 percent of the world's Internet population, but they achieve that goal by having a 100-fold greater number of locations.

The difference in the location count--100:1--is significant. The technologies for next-gen datacenters are designed for locations that generate gigabits of traffic, but Akamai's thousands of small, widely dispersed locations will never serve enough traffic to leverage those economies of scale. Moreover, it's not clear that Akamai's customers--content readers and content providers--get anything other than a marginally better service relative to the datacenter CDNs. The weaknesses in the Akamai approach explain why my vision for a Network Palette--highly-distributed CDNs that add full application-delivery services--has crashed.

So, I believe the centralized approach--i.e., the datacenter model--will wind up winning. However, that said, I don't believe that the players in that space will be the source or inspiration for a new set of applications. Just as no one ever composed a document on a copier, no one will compose applications on this form of successful CDNs and APNs (application platform network). We'll have to look elsewhere for innovations that turn into application platforms.

The Terminal Edge (Total Distribution)

To really reach the edge, you have to go all the way to the end stations. Each process/storage node is dedicated to a single user, and as such, requires

much less security and performance than the shared datacenter solutions.

This edge is actually a programmable application platform. When a browser calls for a Web page, the first thing to arrive is the base-page description. This is actually a small program that tells the browser how to format the page, what elements to get and where to get them. A Web server only needs to send this "program" to the browser, the rest of the content can all come from other servers. This basic programmability has now expanded to java applets, plug-ins and multimedia viewers.

This is the *computing device* vision of the world: computing for the benefit of the user.

Experimentation and user participation can only be found on computing devices at the terminal edge, as was shown via peer-to-peer applications such as Napster and the Grid. However, even peer-to-peer systems use a central service to educate and coordinate the activities of the users. The key condition is that these applications can start small with no additional investment from end-users.

Once these applications become successful, however, economies of scale will pull many of them toward the centralized, datacenter, computing network model. The 10 centers around the globe model is a sure winner for most high-volume applications. Some applications like the Grid will remain operating on end-user devices but will be managed centrally.

Why This Matters

The search for an intelligent edge is actually a search for a sustainable business platform. Many companies are promoting investments in their version of a platform operating at one of the edges to the 'Net. The real question is which approach has the greatest chance for success.

Researchers at Harvard University have studied this question and written, "The Real Options Approach to Standards for Building Network-based Services" by Gaynor, *et al* published at the 2nd IEEE conference on Standardization and Innovation, Boulder, CO, Oct 3, 2001. The authors make a compelling case for the need to have two separate places for applications to reside: In the far edge

under the user's control and in centralized edges under administrative control.

And this isn't an academic issue--lots of folks have a stake in the outcome. First, application developers have to be confident about whether to write applications for user devices directly or indirectly by remote loading of the application and servers. With no other platforms on the horizon, the key question is how to leverage the 100-million desktops that are connected to the 'Net.

Second, CDN customers should investigate alternatives to Akamai. They are likely to find better value, but that'll need to be verified--measure performance all the way from the desktop, not from a measurement point that may be in the rack next to an Akamai server.

Third, enterprises starting to build private CDNs should learn from what's happened on the Internet. Enterprises probably can do very well using modest distribution to regional centers (one per continent is a good start). Focused distribution of content to smaller offices in the off-hours will still make sense, but this is really re-inventing Lotus Notes replication in Web form.

Fourth, many start-ups have been chasing the network edge with interesting products that simply have no market foundation. It's time they got a

grip: It's foolish to try to be the key supplier to the next Akamai when the current Akamai is in trouble.


Finally, there is a lot of work to be done in making datacenters scale. The ISM will be a big market with many players.

Conclusion

We are left with three intelligent edges where money can be made on a sustainable basis: *Network computing* at any datacenter, *computing network* at a system of datacenters and *computing device* in the hands end-users. Of the three, only one is an edge-based new-application platform, the *computing device*. In his keynote address to the eBiz Networks Conference last June Vint Cerf said, "Peter isn't always right." This is one of those times. The Network Palette is dead. Never-used domain for sale, cheap.

Peter Sevcik is president of NetForecast in Andover, MA, and is a leading authority on Internet traffic, performance and technology. He has contributed to the design of more than 100 networks, and led the project that divided the Arpanet into multiple networks in 1984, which was the beginning of today's Internet. He can be reached at peter@netforecast.com.

NetForecast measures application profiles, user behavior and network statistics to predict performance, adoption, and market impact of new technologies. The firm has helped leading service providers, enterprises, and vendors navigate the changing competitive landscape of the Internet economy. We supply key technical and market guidance to ensure the success of network-based projects, products and services.



www.netforecast.com