

To Carriers: Big is Not Beautiful

Net Forecasts – Peter J. Sevcik

BCR Volume 30, Number 11

November 2000

As long as I have been in this business, I have heard the “experts” declare: “Only the big will win.” The only “true” carriers will be global, and offer neatly bundled, multiple, end-to-end services built out of assets they own. Those who own and control the copper and fiber, own and control the network industry.

Moreover, or so I’m told, this is the way customers want it; one simple purchase from a single service provider for all their information needs. Brand and scale are all that matter. The barriers to entry for new players preclude any other outcome.

There certainly is precedent for this line of thinking. Seen any new automobile manufactures recently? How about a new chip fabrication company? The cost of factories and brand development becomes staggering once industries mature.

Since it’s been about 40 years since the first bit went into a modem, and roughly 30 since the first packet popped onto a net, you’d think the industry would’ve reached the “mature phase” by now, one where a few big players take over -- permanently. And some companies -- WorldCom, Qwest and Verizon -- are trying to make that happen through their massive acquisition programs.

Big is Not Beautiful

But, tell me, if the above is true, why is AT&T considering breaking itself up by selling off its consumer long distance telephone service? Why is AT&T Wireless already trading as a tracking stock? The answer has to be that AT&T is not enjoying all the benefits from “synergy” and economies of scale conventional wisdom says it should have.




Figure 1 – High Speed Core Increases Network Hierarchies

And I'm not surprised. I believe the conventional wisdom is wrong; increasingly, bigness produces *diseconomies*. For example, the market research firm, Jupiter Communications, has found that Web sites hit a wall when traffic goes above 1,000 to 5,000 concurrent users. At that point, the cost to support growth becomes non-linear.

Similarly, anyone trying to replicate the U.S. video rental store market would have to build a movie distribution site that, for the peak on Saturday night, could deliver 50,000 *Gbps*, and no one has the faintest idea of how to handle that high a load of simultaneous disk access. It is feasible to replicate a single store on the 'Net, but only if it serves essentially the same size population as the store.

Content delivery networks (CDNs) are often promoted as a way to improve performance, but they actually provide something else: Scale to sites with more content than can be served from a single location. On the 'Net it's actually cheaper to serve a large graphic or video clip from smaller, distributed locations than from one large location. The new approach to network design economies of scale is to break a problem into several layers, some of which may be replicated a thousand times.

The Law of Layering

We already build networks using many protocol layers, because it's more efficient and easier to adopt new technology – the new stuff is slipped into new layers. This is not just an interesting technical way to build networks, it is also how the industry maintains economies of scale. In the networking business, layers are good.

To be sure, end-to-end solutions exist in service industries like networking and transportation. Services need "factories" that will supply each "stage" of the journey in a cost-effective manner. For example, a traveler from Boston to Chicago may engage a taxi (local access) to an airport (transit) to fly an airline (long distance), to another airport and finally a taxi to the destination. The company that is best suited for the job provides each stage of the trip. Frankly, I would be horrified at the prospect of my local taxi company purchasing an airline or *vice versa*. Such an end-to-end service would certainly be less convenient, less safe and more costly than today's layered approach.

The transportation example is a good model for what occurs in network services, because when you boil it down, any network – *every* network – provides three basic functions:

Round 'em up – aggregation of bits/packets/frames/cells.

Set 'em up – switching those bits/packets/frames/cells.

Head 'em out – transmission of...well, you get the idea.

This simple three-step view can be applied to each stage of the process. For example, while the Boston-to-Chicago trip has the three functions operating on a micro level, the three functions reappear at a macro level. The airport provides aggregation via the multiple entry levels. Switching is performed by the walkways and terminal zones (the routing table is on the monitors that identify a gate number) with output buffering performed by the gate lounge. Transmission is single file into the jetway.

What the world wants is good service with alternative choices and competitive prices at each stage of the process. The long-term winning strategy for any network player is to supply a relatively thin portion of the total solution, but at very large scale and extremely low cost. Do one thing very well; resist the temptation to supply the adjacent stages of the journey.

Many Services but Not Many Layers

Now don't get me wrong; I'm not arguing for more fragmentation. No packet should have to traverse more than 3-5 NSPs to get from origin to destination.

However, transport isn't the sole requirement for operating a network. Many other peripheral services are needed -- billing, content delivery, management reporting (see my column on MSPs in *BCR*, May 2000, pp. 8-10), mediation and collection. There are likely to be many new intermediaries that will supply links in the service-delivery value chain but never actually handle a packet.

This flexible approach also fosters new aggregation players with targeted solutions, like Aerzone (www.aerzone.com) (formally Laptop Lane) in airports and InnMedia (www.lodgenet.com) for

hotel guests. These NSPs aggregate services from others and do not build end-to-end solutions. There also will be a need to interconnect local and long distance specialized players with new interchange points. For example, Media Centers (www.mediacentersinc.com) interconnects thousands of retail ISPs and CLECs with the emerging wholesale, optical core, wide-area networks. It is setting up inter-exchange points in large metro areas where these two camps can establish retail-wholesale contracts at data rates of 45 Mbps to 2.5 Gbps.

More importantly, Media Centers has developed a unique OSS (operations and support system) based on XML exchanges between buyer and seller. The contract for bandwidth, special content or user group is negotiated automatically and upon reaching agreement, the connection is established in 14 seconds. Now that's what I call bandwidth on demand, and it will become the *de facto* market for bandwidth and services.

The trap to avoid in this brave new world is having each business layer supplied by many network layers of boxes. Remember that the single largest impediment to good performance on the 'Net is store-and-forward hops (see my column "Performance Issues Facing the World Wide Web" in *BCR*, September 1999, pp. 10-12). The number of hops is climbing as a result of two trends: more business layers and the bandwidth gap between users and the core, a gap that's growing at an alarming rate.

During most of the 1990s, the ratio of the maximum bandwidth in the core divided by the minimum bandwidth at the edge was about 10,000. This means that using the traditional telecommunications hierarchy, equipment with a 4:1 ratio (the telecom world runs on multiples of 4: OC-3, OC-12, OC-48, OC-192) requires 14 boxes to get into the core and back out. By 2005, however, the minimum number of boxes required will jump to 24! Figure 1 shows this dramatic box growth which represents a form of "network friction."

If any speed layer has multiple hops, or if the process has to be repeated for a new business layer (e.g. retail and wholesale) that number could easily double. So, there is a clear imperative: New

networks must be built on equipment with a much greater bandwidth range.

The Network Assembly Line

Lower costs can be achieved by specialization, but we need a new approach -- an assembly-line approach -- to building networks. Henry Ford had a great idea, to build a single car using an assembly line, and he relied on two enabling technologies: Uniform tooling so that parts could be made interchangeable and steel superstructure with which to make large, open buildings.

Similarly, assembly-line networking needs an enabling technology, and fortunately, one's available: Ethernet. In recent months, new specialized players have emerged in local access, metro transit and long haul, and they all rely on Ethernet. It is the standard that permits the interchanging of service layers, and the enabling power of wide bandwidth served from a single node will be available soon (see my column "A Contender With Staying Power" in *BCR*, September 2000, pp. 10-14).

Gotham Networks (www.gothamnetworks.com), for example, is building a new switch that can span many locations, interconnect any port with any service and automate all provisioning functions. The Gotham switch will enable NSPs to build single-purpose solutions on a wide geographic span, and with high port density on a scale that has not been heretofore possible. This assembly-line approach will permit significantly lower costs than the complex multi-service, multi-box, multi-layered approach we've been used to.

Gotham promises a single node that can terminate ports at speed ratios in the tens-of-thousand rather than four. This would permit getting into and out of the core with only four boxes by 2005, and thus foster the spread of business layers while actually reducing the hop count.

Conclusion

The stage is set for a battle between the titans with giant networks and the masses of smaller competitors with "narrow," disruptive networks. My bets on the layered, assembly-line businesses, and I believe they'll become so successful that, some day, we will talk about the monopolistic power

of the winners in some of the key network business layers.

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

