Responses and questions by Bill Woodcock

1. Can IP, ATM or combinations of current protocols scale to the one billion-node Internet?

It's important to think clearly as regards layer-separation, and not fall into all of the fuzzy-thinking traps that led some people into MPLS (too many layer-twos) and IP-over-fiber (too few), or the fuzzy-thinking that led people to hack OSPF into ATM switches for PVC provisioning. The problem at layer three (IP routing, end-to-end connectivity) is fundamentally different than the problem at layer two (adjacencies, meshing, and fast-fail-over).

Both will scale, but the means by which we'll do so are as different as the problems are.

The on-the-wire protocols really aren't a problem, for the most part. It's application-layer stuff like routing decision-making algorithms and stupid content distribution and rate-limiting and a million higher-level things. You can't fix spam, or DoS, or paperwork, with TCP tweaks.
Current routing algorithms incorporate hop-by-hop routing mechanism that makes QoS difficult.
So? QoS was the fad that came between ATM and MPLS. It was tried, found irrelevant, and discarded. The hypothetical world in which it would matter would be one in which pipes were more expensive than people.

What are the requisite properties of new algorithms?

This is obviously the big issue we're all trying to nail down. The fundamental thing that's become obvious since the telco reliability crash forced the issue of dense meshing is that while, price aside, the optimum network topology is obviously a full mesh, our current routing decision-making algorithms were optimized for a _very_sparse mesh, and begin to fail badly as the network itself gets better. Therefore, we desperately need routing decision-making algorithms which work well in an ideal case, as well as the worst-case. Beyond that, propagation and dampening of propagation of policy is the next big requirement. Then authentication of delegations and so forth. I'm sure a lot of people could add to the wish-list.

How can they be tested before fielding?

Ah, _there's_ a really good question, and one I'd love to hear some people propose answers to, or at least lines of inquiry leading to answers.

2. The last mile is a major issue in performance and ability to bridge the digital divide. What kinds of models are needed to address this issue?

This is a huge problem, and entirely a regulatory, political, and business-climate one. Do researchers outside of the economic/political sphere have anything to contribute here?

3. We seem headed toward a ubiquitous Internet characterized as a massively distributed network with powerful local servers forming the master-slave relations. Will this one-level hierarchy be capable of performance and other demands?

This is a misperception based upon inadequate visibility into service-providers' back-end infrastructure. Multiple tiers of infrastructure actually exist behind services like Akamai and Nominum and Digital Island and UltraDNS and even Yahoo and its ilk. Not a problem, and nothing that needs to be addressed in basic research. This is just production engineering, and it's being done adequately.
Is more centralize/hierarchical control over the Internet required?
That's a political and ideological question, not a technological one. It's a topic for debate, not research.

4. Can/should efficient algorithms be designed/implemented that load-balance across geographical servers?

Topological, not geographical. And they've existed for a decade. Again, this is old hat, and production engineering. No research to be done here.
Can/should header size be increased as allowed by future channel capacity increases?
Huh? Cart leading horse? Header size is a pretty basic parameter. Changing it will break a _lot_ of things. You don't change something that basic because it's _possible_, you change it because you _have to_ for some reason. If, at some point in the future, we _have to_ change header lengths, obviously we will. No research to be done here. When we _have to_, it'll be obvious.

Of much more pressing need is universal support for jumbo frames.

> 5. Can the current TCP/IP protocol suite solve the problem of pervasive networking (everything on the internet)?

Sure, no reason why not.
Do we need new architectures that specify new packet designs and the underlying infrastructure?
Nope. New protocols will be defined in response to market demands. The problem isn't creating them, it's stamping them out or making them uniform. How many streaming-audio protocols do we have right now? How many do we need? How do we reduce from the former to the latter?

We don't need to overhaul IP or TCP or UDP... They're basic and functional. Everything interesting has long ago moved up the stack.

> 6. What technologies might push us into next generation of networks? Optical computing and networking, wireless,

Sure, fun stuff. A few things there raise real issues. Wireless has interesting issues in the intersection of regulatory requirements for on-the-air broadcast data and IP encapsulation. The whole MTU issue is starting to get out of hand, and is further aggravated by VLAN-in-VLAN and VPNs and other stuff that people are starting to push in from the edge of the network.

> 7. Are there fundamental problems that must be solved to limit/manage information growth an analogy is energy resources where there exists abundant amount of it, but not practical to harvest it.

Question isn't clear. You want to manage the growth of the Internet by limiting the rate at which people create new information? :-)