Obstacles to Regional HPNV Connectivity And Recommended Action

James O. Whitlock
April 10, 2000

There are now several active plans to deploy broadband (>1Mbps) IP networks among Western New York High Performance Networked Video Initiative (WNY-HPNVI) member groups with anticipation of reliable and seamless High Performance Networked Video (HPNV) connectivity. These include the Regional Community Network (RCN) and the Buffalo Independent Secondary Schools Network (BISSNET) as well as other less publicized plans for new and existing networks. While our experience is certainly limited and we cannot possibly appreciate the priority of HPNV connectivity to the various groups implementing plans, our experience to date suggests several problem areas that could easily preclude seamless connectivity within the region, let alone more ubiquitous global connectivity.

The pattern that we have observed for years is that principals start with narrowly defined objectives within a narrowly defined network or community and then rapidly extend their connectivity objectives outward as applications of the technology are appreciated. This discussion presumes that the initial objective of seamless HPNV connectivity ***within*** the initially defined network domains will not be particularly problematic but that seamless regional (and then global) connectivity will rapidly become both important and seriously problematic if we do not soon start to address the problem areas.

HPNV technologies (and IP videoconferencing in particular) are uniquely affected by impairments in the Internet transport fabric that ***do not affect other network users***. Further, the impairments that are most disruptive are ***not considered serious errors*** within the architectural framework that has guided the development of the Internet since its inception. The Internet is presumed to be an ***unreliable*** transport network that does not guarantee delivery of packets, delivery of packets in sequence, delivery of packets in any particular time (predictable latency), or delivery of packets at any particular rate (predictable bandwidth). As it happens, lost packets, out-of-sequence delivery, jitter (latency variation), and bandwidth congestion (which results in loss, jitter, and sequence errors) critically impair HPNV technologies.

In the long term, the transport fabric will be enhanced to provide committed bandwidth, error free in-sequence packet delivery and bounded jitter for HPNV transport. Work on these areas is proceeding both within the Internet2 community and within the commercial community of network infrastructure suppliers (Cisco, Lucent, Nortel, and others) but many of us believe it will be quite a few years before large public heterogeneous internetworked communities see the results of these developments.

In the meantime, our experience suggests that success with HPNV <u>can</u> be achieved but <u>only</u> when great care is exercised in scaling and tuning the various components in the paths between endpoints. Internet2 sites enjoy a unique status in this regard since the I2 fabric now has substantial headroom (unused available bandwidth) and routers tend to be newer higher performance models. These two conditions drive all the impairments critical to HPNV success towards zero. But our recent experiences with clearing the path to Stanford for the UB/Stanford Bodyworks Class demonstrated that even with I2 endpoints, problems can take substantial time to identify and resolve if for no other reason than that the impairments affecting HPNV are not being routinely monitored because they are not impairments to any other users. TCP, the protocol used for virtually all other data transport on the Internet, corrects and hides the unreliable transport fabric from the end user; UDP, the protocol used by HPNV technologies because of time and performance constraints, does not.

The University at Buffalo contingent has clearly demonstrated, however, that success can be achieved when appropriately knowledgeable attention is directed at critical problem areas. Lisa Stephens' success with the Stanford class and a more recent class taught by Peter Jörgensen from the Netherlands – both of which met or exceeded 384Kbps H.320 quality levels in most regards – further suggest that, within limits, once the dominant impairments have been eliminated by specialist staff for a defined inter-networked community, ad-hoc high quality connectivity can be achieved with normal levels of technical support staffing. While conditions within an inter-networked community will change, they generally do so slowly and the skill-sets and tools necessary to enable initial success should not often be required to sustain it within the same community.

As if network infrastructure and architecture issues were not enough, however, other obstacles related to the very real emergent nature of HPNV technologies will also stand in the way of seamless regional connectivity. These include firewall penetration issues, absence of security and authentication mechanisms, poorly developed management and control infrastructures (Gatekeepers & Neighboring), the generally poor interoperability records of heterogeneous collections of HPNV components, and the paucity of effective diagnostic tools — let alone people to effectively use them with sufficient depth of knowledge in multiple technical domains. Worse still, and regardless of the claims of marketing and sales departments, experienced practitioners regard most HPNV products delivered today as being in a semi-permanent state of beta-test.

I have included below what will undoubtedly be an incomplete list of obstacles to seamless regional HPNV connectivity. While most of the problem areas have been reviewed with the regional network architects, I have seen little evidence to date of clear attempts to address them in network deployment plans. The obstacles include:

1. Insufficient direct connectivity of regional networks.
2. Insufficient local ISP experience with HPNV technologies.
3. Absence of local ISP peering for local traffic.
4. Absence of local success with H.323-aware firewalls.
5. Absence of known methods to achieve secure and authenticated HPNV connections.
6. Absence of a cohesive plans for integrating and Neighboring institutionally managed collections of HPNV components in a controlled, secure, and seamless manner.
7. Minimal amount of interoperability experience with heterogeneous HPNV component collections in mixed I1/I2 environments and generally poor experiences.
8. Minimal regional IT/Network specialist experience with resolving HPNV connectivity and performance problems coupled with a paucity of effective diagnostic tools.
9. Absence of a coordinated shared regional trial and development environment or context.
10. Absence of reliable information on HPNV product interoperability, quality, and reliabiltiy.

Most of these obstacles will not yield to simple solutions but I believe that simple methods can be employed to effectively work towards highly probable success. I offer three suggestions as starting points for discussion and immediate action:

1. Let's do this together and freely share whatever we learn with anyone interested. This has been the cornerstone of our WNY-HPNVI from the beginning. What goes around comes around. To the extent that we share not only with each other – in spite of the competitiveness of some of our parent institutions – but with the larger global community as well, others will share with us. Bob Dixon, Ohio State University and the "Big Ten" Committee for Institutional Cooperation are stellar examples. Let's follow their lead. Moreover, our individual constituencies within the WNY public are highly overlapped and none of us will be able to compete effectively without collaborative solutions to connectivity obstacles. Isolated islands will not serve any of us. Collaboration with competitors, in this regard, is clearly little more than enlightened self-interest.

2. Let's convene a group of technically oriented principals and planners from any and all institutions within the Western New York region who want to achieve seamless regional HPNV connectivity amongst the various existing and new networks. This group could be charged with rapidly assuming command of the issues, developing prioritized and costed recommendations for corrective action, consolidating and analyzing problem reports, collecting and disseminating diagnostic tools and infrastructure configuration experience, and so on.

3. Let's all plan to participate in a large regional trial and testbed that will serve to illuminate, clarify, and prioritize the most serious problem areas. For this purpose, I suggest that the group consider participation in the WNY-HPNVI Extended Grand Rounds (EGR) project now beginning to get underway. More completely defined in a separate document, the EGR project will incorporate live real-time multipoint H.323 videoconferencing, cascaded MCU's, gatewayed endpoints, simultaneous streaming video broadcast in both high performance (Mbps) and low performance (Kbps) modes, capture and presentation for on-demand replay and reinforcement, and possibly a fully produced DVD/stored demonstration multimedia product. As well, the project will embrace WNY endpoints on several or all regional networks of interest, heterogeneous collections of endpoints, multiple management zones behind multiple firewalls, mixed I1/I2 sites, and nationally as well as globally distributed participating institutions. In all

regards, I expect this to be a real-world microcosm and I fully expect a long list of unresolved problems and objectives that are not achieved. I also expect, however, that we will contribute substantially to what is known about lingering problem areas, that we will resolve a number of issues ourselves, and that we will benefit from the interest and contributions of the entire global early adopter community.

I will complete and circulate a description of the WNY-HPNVI Extended Grand Rounds project shortly. It has been in gestation for several years now amongst the more enthusiastic and experienced early adopters in the WNY region but will finally get started within the next few weeks and will continue through the end of the year at least. I encourage you to ignore the medical orientation of the application and to see it as an archetypal HPNV deployment and usage pattern that is equally applicable in all of our spheres of interest. Most of the visions that you have shared with me for applications of HPNV technologies, whether for educational, research, medical, government or business purposes, will employ the same components with minor variations of integration or coupling.

In the meantime, please comment freely and without reservation on the concerns and issues I raise as well as on my straw recommended course of action. If we are to avoid wasting precious time with the scarce resources we have available, we need open candid dialectic and discussion more than polite but insincere agreement or silence. I could easily be misjudging the importance of seamless HPNV connectivity , the magnitude of the issues, or even their true identity. Please speak up and speak freely.