

# NED: a Network-Enabled Digital Video Recorder.

C. Cranor<sup>1</sup>, C. Kalmanek<sup>1</sup>, D. Shur<sup>1</sup>, S. Sibal<sup>1</sup>,  
C. Sreenan<sup>2</sup>, J. van der Merwe<sup>1</sup>

<sup>1</sup>*AT&T Labs- Research  
Florham Park, NJ, USA*

<sup>2</sup>*University College Cork  
Cork, Ireland*

Recent advances in low-cost high-capacity disk drives, and low-cost high-quality video compression and decompression cards are enabling new types of digital video storage and delivery capabilities. One prominent example of this trend is the emergence of Digital (sometimes called Personal) Video Recorders (DVRs) in the home. DVRs enable the digitization, compression, storage, and time-shifted viewing of TV broadcasts. These DVRs typically connect with an analog TV source (e.g., cable TV), interface with an electronic program guide (e.g., via dial-up modem), and are currently capable of local storage of TV content for about 20-80 hours.

Another approach at providing live and on-demand viewing of TV-content is via a server-based Video-on-Demand (VOD) infrastructure. The PRISM project [1], a research effort at AT&T Labs-Research is an example of the latter approach. The PRISM architecture is designed (among other things) to take existing TV content, digitize and compress it in (possibly) network-based servers, and deliver it via an IP infrastructure. In addition to live delivery, network-based storage is used to deliver content shifted in time and on-demand. The PRISM architecture is capable of supporting a number of service scenarios including global access to time-shifted versions of TV content, network based recording, content harvesting based on user specified profiles, as well as more conventional streaming service models such as "movies-on-demand." The PRISM architecture is distinguished from conventional VOD systems by the utilization of a distributed set of network-based servers with storage capabilities, referred to as PRISM portals, to serve streaming content to clients. Clients are connected to a suitable portal using content-aware server selection techniques, and receive the content from the selected portal using conventional streaming protocols such as RTSP and RTP. The Portals are organized into groups of cooperating peers and have the means to discover and extract content from each other. This enables a client to access not only the storage of its "local" portal, but also the content of the entire group of portals. To enable the latter functionality a naming scheme that provides the ability to uniquely identify a particular part of a previously "aired" live stream has been devised. Furthermore, content management mechanisms to efficiently distribute content over the network are employed.

There are many different deployment scenarios for PRISM. For example, Portals can be organized as a "distributed hub" architecture in a single ISP, or as an edge-based architecture (which may span multiple autonomous systems). Portals could be physically located at data centers inside the network, or at the "edge" of the network (such as at cable head-ends or DSLAMs), or even "outside" the network, as long as sufficient bandwidth exists between the portal and the network. It is the content naming, management and discovery protocols that permit this wide range of deployment scenarios, and differentiate PRISM from conventional VOD systems.

A key difference between the home-appliance based and a network-based approach for the delivery of video content is that DVRs are home-based boxes that are not network-enabled, and do not interact with each other or other storage devices. The program storage capacity is limited to that in the DVR itself, and the DVR user

---

<sup>1</sup> Contact: shur@research.att.com

is limited to viewing those hours of programming that s/he had previously selected to be recorded. The PRISM concept involves (potentially vast amounts of) storage in the network, and in principle makes this entire amount of stored content available to the end-user. The drawback of PRISM (and indeed all of all conventional VOD systems) is that the user needs a high bandwidth data connection (e.g., via cable) in order to access the streamed content in real-time. This is because these systems assume that there is no significant amount of local storage, and this assumption places a substantial burden on the QoS capabilities of the network infrastructure. Specifically, the effective throughput for the session between the end-user and the server must be greater than or equal to the rate at which the content is encoded. Packet loss and delay-jitter also needed to be consistently maintained within low limits or else quality degradation may result. (The use of a play-out buffer in clients somewhat alleviates fluctuations in jitter and throughput). If, on the other hand, there is significant local storage on the client, it can be exploited to deliver content much more easily. Specifically, the play-out of content can be de-coupled from the throughput characteristics of the network. Reliable protocols like TCP can be used, which eliminate the problem of packet loss. Content delivery does not have to be in real-time, and can take place ahead of time during intervals in which the network is under-utilized. In fact many of the resource congestion problems that have plagued conventional VOD systems disappear.

In this paper, we describe how to use a network enabled DVR (NED) client to interact with a specific VOD infrastructure, namely, PRISM. Like conventional DVR devices, the NED client allows a user to select specific programs to be stored in the device by means of a program guide. Unlike a conventional DVR device, NED does not limit the choice of content to the set of TV channels that will directly enter the home at some point in the future. Rather, selection can be made from the entire set of PRISM content, including previously stored content, and content that would not normally be delivered to that household by regular means. Unlike a conventional DVR device, the NED client has an IP interface, and executes standard IP-based protocols. The client can request copies of stored TV broadcasts for download (not necessarily in real-time) via the IP interface. This enables the DVR to fetch whatever stored content in the network it has the right to download according to the established content management policy. The use of this client enables the utilization of the PRISM infrastructure, while greatly relaxing the stringent constraints on bandwidth required to receive high-quality streaming video. Because the NED device utilizes an IP interface to download already digitized and compressed video content, it need not necessarily support the encoding or tuner function that conventional DVRs do. A NED device in which this functionality is not required should be less expensive than its non-network-enabled counterpart, because the encoding and tuner hardware can be eliminated. Since the on-demand content does not need to be encoded in real time, more bandwidth efficient off-line coding methods can be used, leading to a better cost-quality tradeoff. In this work we do not address legal matters of copyright, ownership, or security and digital rights management. It is assumed that appropriate legal and business arrangements have been made and digital rights management systems employed, which will allow the download of content in this fashion. In conclusion, we describe the status of the project, including the hardware and software platforms, and briefly discuss next steps.

#### References:

[1] "*PRISM: An IP-based architecture for broadband access to TV and other streaming media*", A. Basso, C. Cranor, R. Gopalakrishnan, M. Green, C. Kalmanek, D. Shur, S. Sibal, C. Sreenan, and J. van der Merwe, Proceedings of the 10<sup>th</sup> Workshop on Network and Operating System Support for Digital Audio and Video, Chapel Hill, NC, June 2000. Available at <http://www.nossdav.org/2000/papers/10.pdf>