




PUBLIC.MEMORY.ORG

Business Plan and Project Description

*Carl Malamud
Senior Fellow/CTO
Center for American Progress
(carl@media.org)*

Table of Contents

<i>Abstract.....</i>	<i>2</i>
<i>Public Proceedings in Washington, D.C.....</i>	<i>2</i>
<i>Two Paths Out of Briefing Rooms.....</i>	<i>2</i>
<i>Congressional Fiber Optic Project.....</i>	<i>3</i>
<i>Overview of Washington Bridge.....</i>	<i>5</i>
<i>Ingesting Data.....</i>	<i>6</i>
<i>Data Center and Peering.....</i>	<i>8</i>
<i>Metadata and Annotation.....</i>	<i>9</i>
<i>Archive Size and Number of Feeds.....</i>	<i>10</i>
<i>Service Goals.....</i>	<i>10</i>
<i>Market Analysis.....</i>	<i>11</i>
<i>Staffing Plan.....</i>	<i>12</i>
<i>Institutional Underpinning and Governance.....</i>	<i>13</i>
<i>The Council of Public Engineers.....</i>	<i>14</i>
<i>Funding Model and Public Partnerships.....</i>	<i>14</i>
<i>Metrics and Milestones.....</i>	<i>15</i>
<i>Timeline.....</i>	<i>16</i>
<i>Pro-Forma Financial Information.....</i>	<i>17</i>
<i>Appendix 1: Sample Router Configurations.....</i>	<i>18</i>



Abstract

The Public Memory Trust (public.memory.org) is a new, nonprofit, nonpartisan public trust incorporated to create public works projects on the Internet and to teach others how to create those public works. The first project of the Trust is the Washington Bridge, a gateway between official briefing rooms in Washington, D.C. and the Internet. The Trust will be incorporated in July, 2006, will commence construction of the bridge in the fall, and will have a beta version working in January, 2007. This business plan begins with the details of the Washington Bridge, then discusses the governance of the trust, the business model for these public works, expected milestones and project timeline, and a pro-forma budget forecast.

Public Proceedings in Washington, D.C.

Spread throughout Washington, D.C. are public briefing and hearing rooms. The White House briefing room and congressional hearing rooms are perhaps the most visible, but most agencies maintain such facilities for briefing the press or holding public hearings or other proceedings. Some of these events are covered for television or the Internet by major media such as C-SPAN (www.c-span.org) or by specialized trade services such as FedNet (www.fednet.net). Some agencies and congressional committees provide their own “webcasts” of audio and video from their proceedings as well.

While modest progress has been made in making audio, video, and other data from official events available on the Internet, there is no systematic effort to provide comprehensive coverage and there is no permanent archive available. To observe the official events of our government, one must be physically present in Washington, D.C. at the time of the event.

Two Paths Out of Briefing Rooms

When an agency or a committee provides coverage of their own event, the current “state of the art” usually consists of shanghaiing the person responsible for computer support, handing him or her a cheap camera, and hiring a webcasting company to send the data out to the Internet. If an event is popular, many streams are created, and the committee or agency bears a significant per-user financial cost. Archives are maintained on an ad-hoc basis by each group. In short, there is a significant amount of work and expense associated with providing access to the public over the Internet, the result being no comprehensive archive and no systematic coverage of events.

A second pathway out of an official briefing uses the media as a gateway. Most hearing rooms have significant support for the media, including a sound system that supports a “mult box,” an audio splitter that provides feeds from all the speaker microphones to any media that are present (see figure 1) . A reporter, for example, will often plug a tape recorder into the mult box as a way of creating an audio recording of the proceeding for future reference.



Figure 1. The VA-32 “Press Feed System Mult Box” from OPAMP Labs

Television crews also use these mult boxes, but usually acquire video through their own cameras. In some cases, the crews record straight to tape, but often there is a requirement to send the video in real-time back to production facilities. For this, most official proceedings in Washington use the services of the Audio Visual Operations Center (AVOC), a service provided by Verizon Communications.

The AVOC is a video “switch” that uses fiber optics to transmit broadcast-quality video streams. Media organizations subscribe to the AVOC service by leasing a switch, and dedicated fiber optic runs to hearing rooms and back to their production facilities. When a particular hearing is being covered, the camera crew connects their camera to their fiber outlet in the hearing room. The production facility then issues a software command to the switch in the AVOC that provides a connection between the run to the hearing room and one of the runs back to the production facility. Figure 2 shows a high-level diagram of this process.

Congressional Fiber Optic Project

Most federal agencies have one or a small number of locations from which media crews are expected to film. For these agencies, Verizon has direct connections through the building wiring into the briefing rooms. On Capitol Hill, however, there are a large number of hearing rooms and a second layer of infrastructure is used.

Media access to Capitol Hill is controlled by the Radio-TV galleries. The galleries certify broadcast journalists and their sponsoring institutions for access. One of the long-term projects of the Radio-TV galleries has been a Fiber Optic Project which now connects a large number of hearing rooms. Through financial participation in the Fiber Optic Project, gallery members may access this infrastructure which provides distribution from the Capitol Complex to sub-hub rooms in the Dirksen Senate Office Building and the Rayburn House Office Building (see figure 3). Those sub-hubs in turn connect to a large number of (eventually all) hearing rooms.

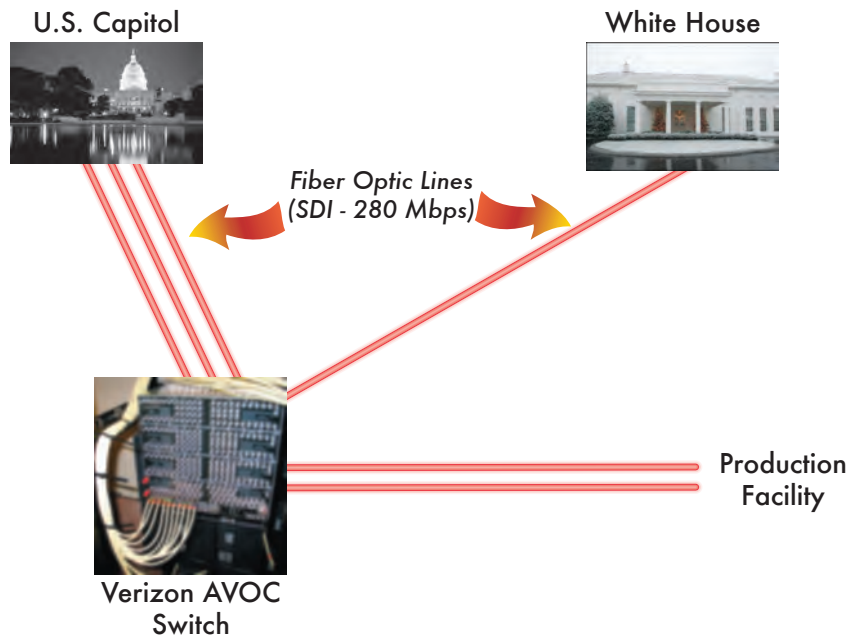


Figure 2. Overview of Verizon's AVOC Service



Figure 3. Fiber Optic Project in Capitol Complex

Broadcasting from Capitol Hill thus involves two stages of infrastructure. First, gallery members connect a particular hearing room to their equipment in the hub room. Then, a connection is made to the member's AVOC lines for transportation back to the AVOC switch and finally on to the member's production facilities.

Overview of Washington Bridge

The Washington Bridge is a public works project that aims to connect a large number of hearing rooms in Washington, D.C. to the Internet. An overview of the bridge is shown in figure 4.




Figure 4. Overview of the Washington Bridge

There are three parts to the bridge, each of which will be discussed in turn:

1. Connecting hearing rooms to the AVOC. Initial plans are for 8 AVOC lines to the Congress, 8 lines to agencies, and 16 lines from the AVOC back to the control room.
2. Taking data from the AVOC, which is video over fiber, and turning those feeds into Internet feeds occurs in the Control Room, located close to the AVOC. Encoded video is then transported over the Internet Protocol (IP) on a high-speed line to a data center installed in a co-location facility.
3. Each of the feeds is then processed in a variety of ways, including signing the feeds and transforming them into various lower-resolution encodings, where they are then made available at no charge to commercial networks and research networks. The bridge will use direct peering with many networks at the co-location facility, will

create a direct interconnection into [Internet2](#) and [LambdaRail](#) (the two major research networks), and will also create transit agreements to provide access to networks not directly located at the peering points.

Ingesting Data

Figure 5 shows an overview of the process of obtaining video feeds.

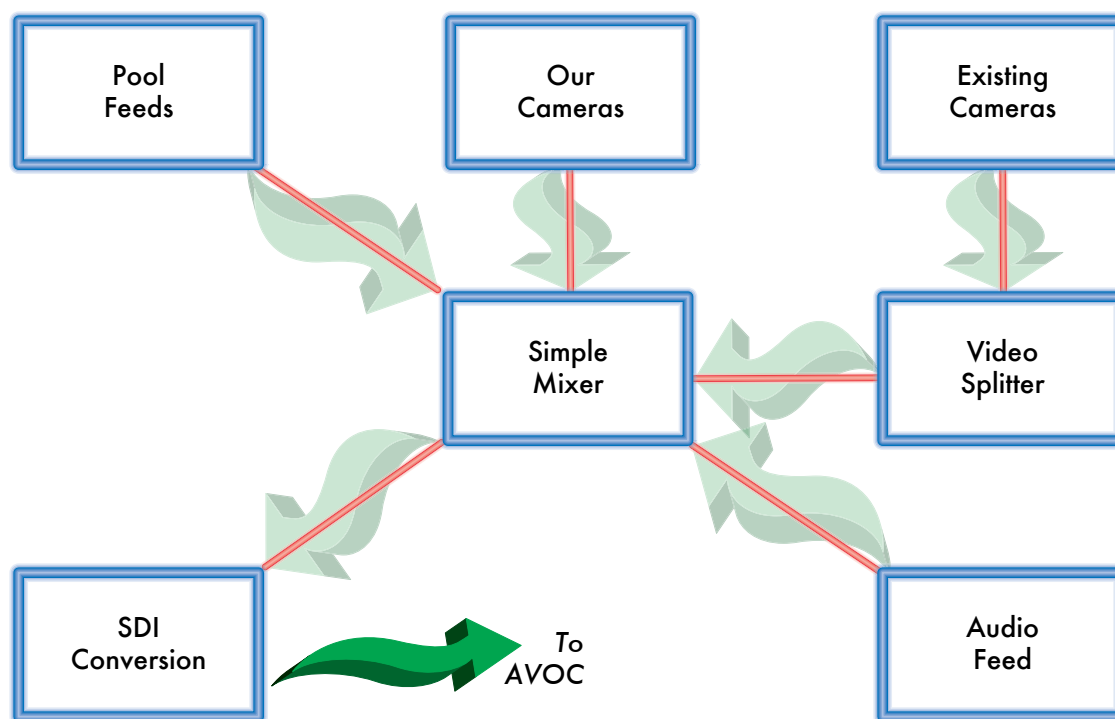


Figure 5. Ingesting Video Feeds

There are three sources of feeds:

1. In many cases, coverage of events is limited to a small number of cameras. Members of the various Correspondent's Galleries often engage in pooling arrangements, where one member provides the cameras and other members a pool copy of the video. Participation in pools is a matter of reciprocity: a member must provide the others with feeds on occasion.
2. As members of a Correspondent's Gallery or through other permission from the relevant agency, public.memory.org can bring in cameras and directly film the event. We anticipate working with "prosumer" gear, such as \$3,000 3-chip cameras.
3. Some agencies already have cameras. The video quality is not the best, but leveraging existing feeds is a quick way to add content through the simple mechanism of splitting the video signal. In the long run, it is hoped that there will be incentive for existing agencies to upgrade the quality of their audio-visual equipment.

Obtaining audio is usually a simple matter of connecting to an existing “mult box.” The audio and video are mixed together if necessary, the signal is converted to the Serial Digital Interface (SDI) standard, and the feed goes back to the control facility.

The Washington Bridge will have a capacity of 16 total lines of video fed from the AVOC switch (see Figure 6). At the control room, which will be located in downtown Washington, D.C. close to the AVOC, standard cable head-end equipment will be used to turn the data from 280 Mbps video-over-fiber SDI feeds into MPEG-compressed video transported on the Internet Protocol (IP).¹




Figure 6. High-Level Diagram of Control Room Components

There are a variety of ways to convert SDI into MPEG. Two examples are encoders manufactured by Scientific Atlanta. The first uses the DVP series encoders to create a 50 Mbps MPEG-2 encoded video stream. The second uses the DCM series encoders, which offer more flexibility in output formats, such as MPEG-4 encoding.

These video feeds, subsequent to being encoded, are then input into a router, and sent over an OC48 Internet line, which runs at 2.4 billion bits per second (Gbps).² The feeds are routed to the main data center. The philosophy of design is to turn the dedicated video-over-fiber into IP-based data as close to the source as possible. Once turned into

¹ The Serial Digital Interface (SDI) is defined by the Society of Motion Picture and Television Engineers (SMPTE) in standard SMPTE-259M-1997 and by the International Telecommunication Union in [ITU-R BT.656-4 \(02/98\)](#). The Moving Picture Experts Group (MPEG) has defined [MPEG-2](#), which is widely used in products such as digital television set top boxes and DVDs, [MPEG-4](#), which is being increasingly used for web-based delivery, as well as some emerging standards such as [MPEG-7](#), a standard for description and search of content in video.

² The Optical Carrier (OC) levels are a series of service levels in the Synchronous Optical Network (SONET) framework defined by the American National Standards Institute (ANSI) in standard [ANSI T1.105-2001](#).

standard Internet traffic, the remaining process can be done in a high-quality data center which has power, security, and other protections necessary for the bulk of the equipment.

Note that final decisions on control room location and data center location will occur during the summer engineering process (see [Timeline](#)). The MPEG-encoding of the data may occur in the facilities of the Public Memory Trust, or may be located inside of a data center located close to the AVOC.

Data Center and Peering

The bulk of the data processing occurs at the data center. A likely location would be the PAIX co-location facility run by Switch and Data. Figure 7 shows a high-level diagram of equipment anticipated to be installed at the data center.

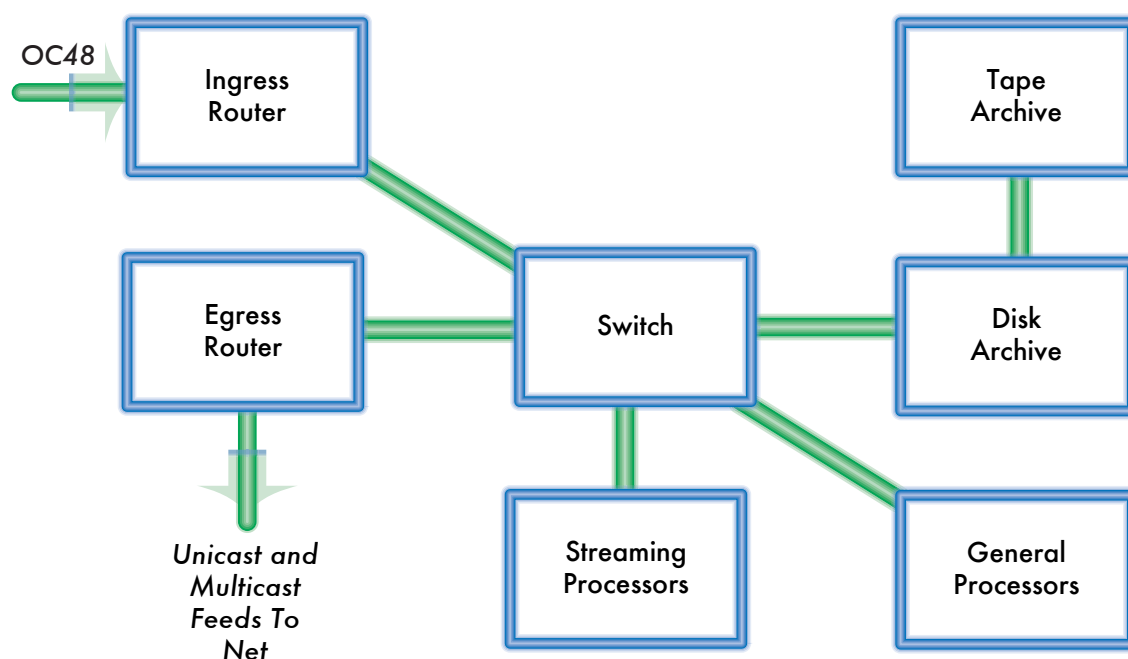


Figure 7. High-Level Diagram of Data Center Components

Conceptually, the data come in from the control facility as 16 video feeds embedded in a direct IP connection. These feeds are received at the ingress router shown in Figure 6, which is responsible for accepting the incoming 16 feeds from the Control Room near the AVOC. The ingress router, in turn, connects to a high speed switch which makes the data available to a variety of general purpose processors:

1. A high-volume disk system (e.g., 1 petabyte, which is equal to 1 million gigabytes), which is used to create the initial permanent archive. The permanent archive is in turn connected to a tape jukebox, which allows the creation of replicas of the archive. These replicas can be transmitted to other institutions, such as the nearly 1,250 libraries that participate in the [Federal Document Repository Program](#).

2. In addition to the creation of the permanent archive, another set of general-purpose processors will be used to perform a variety of transformations on the feeds. These transformations include the generation of metadata through automated methods, transcoding into lower-resolution formats, and signing operations to be used for future data authentication.
3. The third set of processors are used to support streaming of data, either live feeds or from the archive, to end users. As will be discussed subsequently, the Public Memory Trust will maintain a limited end-user presence which provides direct access to the data. In addition, a goal of the Public Memory Trust will be to feed this video to a variety of places, including commercial and non-commercial operations, so they can in turn serve their end-users.

The main interface to the rest of the Internet is an egress router, which is a high-capacity system able to provide direct peering at high performance to a large number of networks. Engineering for the data center configuration will be finalized during the summer (see [Milestones](#)), however [Appendix 1](#) contains sample configurations for the ingress router, the switch, and the egress router.

Metadata and Annotation

One of the most challenging aspects of providing the Washington Bridge service is generating metadata, which can be used to identify events and search for particular subsets of events. The traditional approach to generating this metadata is manual: human beings providing functions such as identification of keywords, generation of verbatim transcripts, and other “value added” information.

For this service, we intend to pursue three approaches for the generation of metadata:

1. A minimal manual approach, in which operational staff provide “bootstrap” information that is readily available. Because staffing is limited, this bootstrap information consists of data such as the time and date, location, and brief descriptive information. It may also include information generated through the scraping of existing web pages provided by the agency, or the collection and scanning of papers distributed at the event.
2. To the extent possible, the service will use as many automated techniques as can be found. A variety of software tools provide functions such as speaker segmentation, speaker identification, scene identification, and rudimentary speech to text identification.
3. One of the most powerful aspects of the Internet is the ability to draw on a community of volunteers to add value to public domain data. As such, we intend to undertake a prototype effort to couple instant messaging technologies which are

based on the XMPP standard³ to archives and relevant streaming solutions. Because instant messaging uses timestamps, as do the various video formats, we believe it will be possible to provide a service where “anybody can annotate” a video stream. Annotation can include commentary, transcriptions, and pointers to additional information. A user-based rating and tagging system can be used to make the wheat rise from the chaff.

Archive Size and Number of Feeds

The size of the archive depends on three factors: the encoding format chosen for archiving, the number of corollary streams associated with the basic footage, and the number of hours. For this calculation, we assume 50 Mbps MPEG-2 streams and additional data of 25 Mbps which includes alternative encodings. We assume 16 AVOC streams with 5 hours per day each of content. This yields an archival rate of 0.675 petabytes/year. Through the use of various optimizations (e.g., selection of 8 Mbps MPEG-4 as the archival format), the total size of the archive needed is estimated approximately 1 Petabyte to cover the 3-5 year period of operations.

For planning purposes, we have sized this service based on a total of 16 AVOC lines, 8 to the Capitol and 8 to agencies. The numbers of lines to the Capitol is based on peak activity during periods of high congressional activity. For example, on the week of May 15, 2006, in addition to active the floors of the House and Senate, there were typically simultaneous events in 7 Senate hearings that started in the morning and 1-2 additional events that started in the afternoon. During this week, there happens to be few House hearings scheduled. As can be seen, the amount of programming available is highly cyclical. A capacity of 8 lines as a planning target seems like an appropriate compromise between installing enough capacity to handle peak and the issue that many of the lines will sit idle during periods of congressional inactivity.

Service Goals

All services provided from the Washington Bridge attempt to balance two potentially competing objectives:

1. To show that the service is real, a “retail” presence has to be provided. By “retail,” we mean direct service to the end user.
2. The eventual goal of the Washington Bridge is to see this service provided by a distributed collection of players. In particular, we would like to see commercial providers (e.g., Google Video and other services on the net), major media intermediaries, and research/educational providers (e.g., university libraries that are part of the Federal Document Depository Program) providing direct service to the end user.

³ The Extensible Messaging and Presence Protocol (XMPP) is a set of instant messaging standards defined by the Internet Engineering Task Force (IETF) in RFCs [3920](#), [3921](#), [3922](#), and [3923](#). Extensions to XMPP are standardized by the Jabber Software Foundation in the [JEP Series](#).

Note that the terms “wholesale” and “retail” do not imply a charging model: the services will all be provided at no charge.

The goals are somewhat conflicting. For the retail service, it is important to add as much value as possible to the data in order to make the service useful to the end user. On the other hand, the better the service is, the more it “competes” with the services created by the wholesalers who ingest the full data feed.

We expect the staff of the Public Memory Trust, with the guidance of the board, to balance these two competing goals. By devoting some energy to providing a good end-user presence, the service helps establish a standard of “minimal acceptable functionality” which can easily be provided. On the other hand, by encouraging wholesalers to ingest the data feed and create their own services, we meet the longer-term goal of spreading this service into the underlying Internet infrastructure.

An important goal of the Washington Bridge service is to function as a research testbed that allows partners of the Public Memory Trust to try out new techniques and ideas. Examples of such efforts include:

- Signing video streams using a signed checksum, such as MD5 or SHA1.⁴
- Speech-to-Text generation of metadata for searching.

The ultimate success metric for the creation of a service such as the Washington Bridge is to put ourselves out of business because that infrastructure is being created in a distributed fashion by many players.

The model used for the Washington Bridge applies to many other scenarios, including state & local governments in the United States and all levels of government in other countries. The long-term goal of the Public Memory Trust is to replicate our knowledge.

Market Analysis

Is there a compelling need for video from official hearings on the Internet? Will people really care? Answering that question is difficult because the service described here does not exist. However, two trends in traffic on the Internet make a fairly convincing case that this data would be well received.

First, video on the Internet has exploded. According to the Associated Press, [YouTube.Com](#) was publishing 35,000 new videos per day in February, 2006 and users were viewing 35 million videos per day. The site was reaching an audience of 9 million U.S. visitors per month and ranked 31 on [Alexa.Com](#)’s traffic rankings. YouTube is one of several dozen sites that provide an outlet for video, a market that includes large distributors such as [Google Video](#), Apple’s [iTunes Music Store](#), and a large number of open-source or community-based services.

⁴ Message Digest 5 (MD5) and Secure Hash Algorithm 1 (SHA1) are cryptographic algorithms used to ensure message or file integrity and were standardized by the IETF in RFCs [1321](#) and [3174](#).

The second trend is the immense popularity of bloggers who write about current events. [Technorati.Com](#) tracked over 39.4 million blogs for April, 2006, some of which are getting close to the size of “major media.” For example, [BoingBoing.Net](#) had links into it from 20,223 other web sites, and is reaching audiences of 2.5 million viewers per month (see [boingboing.net/stats/](#) for current statistics).

Most end users will not view an entire congressional hearing end-to-end. The audience for a full event will probably consist of scholars, students, Washington insiders, the media, and the bloggers. In many cases, however, we expect edited versions or excerpts of events to reach very large audiences. Both major media and the growing number of commentators and video producers on the Internet will use the raw material to incorporate into their publications.

Staffing Plan

The staffing plan for this service is based on 12 full-time positions as shown in figure 8:




Figure 8. Organization Chart

Staffing is based on the assumption that core staff will be supplemented by the [Council of Public Engineers](#), the advisory body. Salaries will be kept low, but because of the scope and visibility of the project and the distinguished nature of the advisory body, we hope to attract top-caliber mid-career professionals anxious to work on a project of large scale as a stepping-stone in their careers.

The organizational chart consists of 5 Internet specialists, 5 AV specialists, and a facilities manager. We anticipate all staff being able to operate cameras, and thus provide the ability to produce a large number of on-site events during peak periods. We also expect

to supplement the permanent staff with interns, fellows (professionals whose salary is paid by their employers), and volunteers who have a research interest in the work.

Institutional Underpinning and Governance

The institutional underpinning for the construction of the Washington Bridge is the Public Memory Trust, a 501(c)(3) nonprofit trust with a charter to create public works projects on the Internet and teach others how to create public works projects.

Governance of the Public Memory Trust will be through the selection of a nonpartisan Board of Directors. The initial Board of Directors will be selected upon incorporation of the Trust, and will serve terms of 2 or 3 years to provide staggered succession. The Executive Director will report to and work under the direction of this Board of Directors.

We anticipate an initial board of 13 members, selected through 3 mechanisms:

1. Two directors appointed at incorporation.
2. Five directors, one named by each of five founding sponsors.
3. Five “outside” directors, named by the first 8 selected. Outside directors should be prominent figures from media, government, or technology that will be selected on the basis of their knowledge and their ability to represent the overall public interest.

Selection of board members after initial terms expire will be based on a self-perpetuating board model, where the current members elect their successors. However, with a supra-majority vote, the board may alter the articles of incorporation to move towards a member-based governance mechanism, such as having the Council of Public Engineers select board members.

Governance of the Public Memory Trust will be vested in the Board of Directors. Several principles will be written into the articles of incorporation:

1. Full transparency shall be an explicit goal of the Trust, including full disclosure of all financial information, and explicit specification of peering and coverage criteria.
2. The Board shall not be directly involved in programming decisions or determination of how data will be distributed on the net. Instead, the Board shall require the staff, in consultation with the Council, to develop explicit criteria on how such decisions are being made. The Board shall then review and approve those criteria to ensure that they are objective and neutral.
3. To provide full oversight over operations, the Board shall appoint an executive committee. If the board consists of 13 members, an executive committee of 5 would provide continuity.

The Council of Public Engineers

Technical advice and high-level design principles for the Washington Bridge are provided by a technical advisory group, known as the Council of Public Engineers. This group will operate on the rough consensus model which forms the basis of other standards-making bodies, such as the Internet Engineering Task Force.⁵

The initial composition of the Council consists of 3 people named by each of the founding sponsors and the incorporator. Members should be named based on their technical qualifications. It is expected that the Council will formally meet once per year and that other activities shall be conducted via the Internet through a web site and/or WIKI, a bulletin board, and electronic mailing lists. Because all members of the Council will be extremely busy people, great pains will be taken to use their time wisely.

Funding Model and Public Partnerships

Initial funding for the Public Memory Trust is based on a dedicated fund sufficient to create and operate the Washington Bridge. Funding is anticipated to come in three forms:

1. In-kind contributions of labor through a fellows program, and by volunteers and advisors.
2. In-kind contributions of equipment.
3. Cash contributions.

The over-all cash budget for the Trust is approximately \$2.5 million/year, assuming the bulk of capital expenditures are covered through in-kind contributions for routers and servers. After factoring in capital expenditures for items not covered by in-kind contributions as well as startup-costs, cash requirements for the period through the end of 2009 are \$8 million. The cash budget is thus based on 4 sponsors at the \$2m contribution level, with the possible reduction in cash expenses through the fellows program and other in-kind contributions. The budget assumes in-kind contributions of hardware of approximately \$4 million.

A long-term goal is dispersion of our work into the fabric of Washington, D.C. Ultimately, this responsibility will fall on the individual agencies, the Congress, and the Government Printing Office (GPO). To facilitate this long-term goal, we will attempt to negotiate partnership agreements with relevant agencies, such as GPO, in which Federal employees can be detailed to the Public Memory Trust as Fellows, enabling them to learn from the on-going operations. In addition, we hope to establish a number of other partnerships with non-profit organizations, such as community media centers and Internet-based organizations such as [the Internet Archive](#).

⁵ The basic definition of the IETF standards process is in [RFC 2026 \(BCP 9\)](#).

Metrics and Milestones

Because the Washington Bridge provides a new service, it is hard to predict how many users the service will have. To aide in future evaluations, we have set out 4 broad sets of metrics and milestones with which to evaluate the service:

1. A public beta which provides direct service to end users will be up and running in Q1 of 2007. We expect a wide range of usage patterns for different events. Based on popular video services such as YouTube, we expect usage to vary from a low of a cumulative number of total users of 30,000 viewers up to perhaps 3 million viewers for a few more popular events.
2. The service is based on the wholesale/retail model. A success metric is to have, by the end of 2008, at least two large commercial providers ingesting data feeds and providing a value-added service as well as at least two large educational institutions, such as libraries, providing permanent archives.
3. The long-term goal of building the Washington Bridge is to see this service provided on a permanent basis. In an ideal success scenario, by 2010 the Government Printing Office will have taken over the core switching functions, individual agencies and Congressional committees will all be providing high-quality video directly out of their facilities, and a large number of members of the Federal Document Repository Program will have undertaken the provision of a permanent archive.
4. By 2010, the mission of replicating the service to other locations should be well underway. A goal would be having the Public Memory Trust involved in the design and creation of at least two other services, and the provision of consulting and education for at least 20 other locations. A corollary goal is a long-term funding model for the provision of these subsequent projects.

Timeline

As shown in Table 1, the summer of 2006 will be used to finalize sponsor agreements, incorporate the new entity, and finalize the initial engineering. The service will be built in Q4/2006, will enter a public beta period in Q1/2007, and will be in full production in 2008. A new set of features for version 2.0 will be introduced and stabilized in 2009, and in 2010 we will find a long-term institutional framework for the Washington Bridge.

Date	Goal
5/2006	<ul style="list-style-type: none"> • Finish technical briefings to potential sponsors.
6/2006	<ul style="list-style-type: none"> • Initial sponsor agreements reached.
7/2006	<ul style="list-style-type: none"> • Incorporation of Public Memory Trust. • First hires.
8-9/2006	<ul style="list-style-type: none"> • Finalize engineering plans. • Placement of orders for fiber and equipment.
Q4/2006	<ul style="list-style-type: none"> • Creation of basic service.
Q1/2007	<ul style="list-style-type: none"> • Service is operational in beta mode.
Q2-Q4/2007	<ul style="list-style-type: none"> • Continued development of service.
2008	<ul style="list-style-type: none"> • Service is fully operational.
2009	<ul style="list-style-type: none"> • Additional functionality added.
2010	<ul style="list-style-type: none"> • Transition of service to permanent infrastructure providers.

Table 1: Project Timeline

Pro-Forma Financial Information

	2006	2007	2008	2009
STAFF EXPENSE				
Payroll	336,000	1,120,000	1,187,200	1,258,432
Contractors	40,000	80,000	80,000	80,000
Benefits & Taxes	67,200	224,000	237,440	251,686
Travel	30,000	50,000	50,000	50,000
OFFICE EXPENSE				
Rent	53,333	160,000	160,000	160,000
Insurance	12,000	12,000	12,000	12,000
Legal	5,000	5,000	5,000	5,000
Accounting	7,500	12,000	13,000	13,000
BANDWIDTH				
AVOC	18,000	72,000	72,000	72,000
Transit	6,000	60,000	120,000	120,000
MAX	12,000	130,000	130,000	130,000
PAIX	18,000	72,000	72,000	72,000
OC48	33,000	132,000	132,000	132,000
OTHER				
Tapes & Shipping				
Pool Feeds	10,000	30,000	30,000	30,000
CAPITAL EXPENSE				
Video Equipment	50,000	50,000		
Control Room	100,000	50,000		
Data Center	100,000			
TOTAL	898,033	2,259,000	2,300,640	2,386,118

Appendix 1: Sample Router Configurations

Product	Description	Q	Price	Total
CISCO7604	Cisco 7604 Chassis Bundle	1	\$0	\$0
FAN-MOD-4HS	High-Speed Fan Module for 7604/6504-E	1	\$0	\$0
7604-SUP7203B-PS	Cisco 7604 Chassis, 4-slot, SUP720-3B, PS	1	\$34,000	\$34,000
SUP720-3B	Supervisor Engine 720-3B	1	\$0	\$0
MEM-C6K-CPTFL256M	Cat6500 Sup720/Sup32 Compact Flash Mem 256MB	1	\$900	\$900
MEM-S2-512MB	Catalyst 6500 512MB DRAM on the Supervisor (SUP2 or SUP720)	1	\$3,600	\$3,600
MEM-MSFC2-512MB	Catalyst 6500 512MB DRAM on the MSFC2 or SUP720 MSFC3	1	\$3,600	\$3,600
2700W-AC	2700 W AC Power Supply for 7604	1	\$0	\$0
GLC-T	1000BASE-T SFP	1	\$395	\$395
CAB-GSR16-US	Cisco 12016 GSR AC Power Supply Cord, US	1	\$0	\$0
S763ISK9-12218SXF	Cisco 7600-SUP720 IOS IP SERVICES SSH	1	\$5,000	\$5,000
OSM-20C48/1DPT-SS	2-port OC-48/STM-16 POS/DPT OSM, SM-SR, with 4 GE	1	\$90,000	\$90,000
MEM-OSM-256M	256 MB ECC Memory for Optical Services Modules	1	\$7,500	\$7,500
WS-X6748-GE-TX	Cat6500 48-port 10/100/1000 GE Mod: fabric enabled, RJ-45	1	\$15,000	\$15,000
MEM-XCEF720-512M	Cat 6500 512MB DDR, xCEF720 (67xx interface, DFC3A/DFC3B)	1	\$4,800	\$4,800
WS-F6700-DFC3B	Catalyst 6500 Dist Fwd Card, 256K Routes for WS-X67xx	1	\$7,500	\$7,500
WS-SUP720-3B	Catalyst 6500/Cisco 7600 Supervisor 720 Fabric MSFC3 PFC3B	1	\$0	\$0
PWR-2700-AC/4	2700W AC Power Supply for Cisco 7604/6504-E	1	\$0	\$0
CON-SNT-SCSO7604	SMARTNET 8X5XNBD Top Lvl-Svc on each component	1	\$0	\$0
CON-SNT-7604BP5	SMARTNET 8X5XNBD 7604 Chass, 4slot SUP720-3B, PS	1	\$2,720	\$2,720
CON-SNT-20C48D5S	8x5xNBD Svc, 2-prt OC48/STM16 POS/DPT OSM,SM-SR	1	\$10,400	\$10,400
				\$185,415

Table A1-1: Sample Configuration of Ingress Router at Data Center

Product	Description	Q	Price	Total
WS-C6503E-S32-10GE	Cat6503E chassis, WS-SUP32-10GE-3B, Fan Tray (req.P/S)	1	\$22,995	\$22,995
SC6K-S323K9-8.5	Catalyst 6000 Sup 32 PFC3 Flash Image w/SSH, Rel 8.5	1	\$0	\$0
MEM-XCEF720-512M	Cat 6500 512MB DDR, xCEF720 (67xx interface, DFC3A/DFC3B)	1	\$4,800	\$4,800
MEM-C6K-CPTFL256M	Cat6500 Sup720/Sup32 Compact Flash Mem 256MB	1	\$900	\$900
XENPAK-10GB-CX4	Cisco 10GBASE-CX4 XENPAK Module	1	\$600	\$600
WS-X6148-GE-TX	Catalyst 6500 48-port 10/100/1000 GE Mod., RJ-45	1	\$7,500	\$7,500
PEM-20A-AC+	PwrEntryMod use w/1400W AC P/S for CISCO7603, WS-C6503	1	\$250	\$250
PWR-1400-AC	1400W AC pwr/sup for CISCO7603 and Catalyst WS-C6503 chassis	1	\$745	\$745
CAB-AC-C6K-TWLK	Power Cord, 250Vac 16A, twist lock NEMA L6-20 plug, US	1	\$0	\$0
WS-SUP32-10GE-3B	Cat 6500 Supervisor 32 with 2 ports 10GbE and PFC3B	1	\$0	\$0
SM2AIPBK9-122175XB	Cisco CAT6000-MSFC2A IOS IP BASE SSH	1	\$0	\$0
MEM-MSFC2-256MB	Catalyst 6500 256MB DRAM on the MSFC2	1	\$0	\$0
WS-C6503-E-FAN	Catalyst 6503-E Chassis Fan Tray	1	\$0	\$0
CON-SNT-6503E10G	SMARTNET 8X5XNBD Cat6503E, WS-SUP32-10GE-3B Fan Tray	1	\$1,680	\$1,680
				\$39,470

Table A1-2: Sample Configuration of Main Switch at Data Center

Product	Description	Q	Price	Total
GSR10/200-AC	Cisco 12410 200 Gbps; 1PRP, 2 CSC, 5 SFC, 2 Alarm, 2 AC	1	\$95,000	\$95,000
PRP-2	Cisco 12000 Performance Router Processor 2 (PRP-2)	1	\$16,000	\$16,000
MEM-PRP2-1G	Cisco 12000 PRP-2 1Gig DRAM Default Option	1	\$0	\$0
MEM-12KRP-FD1G	Cisco 12000 Series, 1GB PC ATA Flash Disk	1	\$8,900	\$8,900
HD-PRP2-40G	Cisco 12000 PRP-2 40G Hard Drive	1	\$5,000	\$5,000
CAB-GSR16-US	Cisco 12016 GSR AC Power Supply Cord, US	2	\$0	\$0
S12K52-12.0.32S	Cisco 12K Series PRP IOS SERVICE PROVIDER/SECURED SHELL 3DES	1	\$13,500	\$13,500
12000-SIP-601	Multirate 10G IP Services Engines (Modular)	1	\$99,000	\$99,000
SPA-1XTENGE-XFP	1-port 10Gigabit Ethernet Shared Port Adapter XFP based	1	\$20,000	\$20,000
XFP-10GLR-OC192SR	Multirate XFP module for 10GBASE-LR and OC192 SR-1	1	\$4,000	\$4,000
12000-SIP-601	Multirate 10G IP Services Engines (Modular)	1	\$99,000	\$99,000
SPA-10X1GE	10-port Gigabit Ethernet Shared Port Adapter	1	\$30,000	\$30,000
SFP-GE-S	1000BASE-SX SFP (DOM)	10	\$5,500	\$55,000
12000-SIP-601	Multirate 10G IP Services Engines (Modular)	1	\$99,000	\$99,000
SPA-2XOC48POS/RPR	2-port OC48/STM16 POS/RPR Shared Port Adapters	2	\$152,000	\$304,000
SFP-OC48-SR	OC-48c/STM-16c SFP, Short Reach	4	\$10,000	\$40,000
12000-SIP-601	Multirate 10G IP Services Engines (Modular)	1	\$99,000	\$99,000
SPA-OC192POS-XFP	1-port OC192/STM64 POS/RPR XFP Optics	1	\$125,000	\$125,000
XFP-10GLR-OC192SR	Multirate XFP module for 10GBASE-LR and OC192 SR-1	1	\$4,000	\$4,000
12000-SPA	SPA for Cisco 12000; No Physical Part; For Tracking Only	1	\$0	\$0
12000-SPA	SPA for Cisco 12000; No Physical Part; For Tracking Only	1	\$0	\$0
12000-SPA	SPA for Cisco 12000; No Physical Part; For Tracking Only	2	\$0	\$0
12000-SPA	SPA for Cisco 12000; No Physical Part; For Tracking Only	1	\$0	\$0
CON-SNT-GSR10/200	8x5xNBD Svc, GSR 12010 & 12410 200 (GSR10/200)	1	\$9,600	\$9,600
CON-SNT-PRP2	SMARTNET 8X5XNBD Cisco 12000 Performa	1	\$2,240	\$2,240
CON-SNT-1XTENXFP	SMARTNET 8X5XNBD 1-port 10Gigabit Ethernet Shared Port	1	\$800	\$800
CON-SNT-10X1GE	SMARTNET 8X5XNBD 10-port Gigabit Ethernet Shared Port	1	\$1,200	\$1,200
CON-SNT-2XOC48PO	SMARTNET 8X5XNBD 2-port OC48/STM16 PO	2	\$6,080	\$12,160
CON-SNT-SIP601	SMARTNET 8X5XNBD Multirate 10G IP Ser	1	\$3,960	\$3,960
CON-SNT-SIP601	SMARTNET 8X5XNBD Multirate 10G IP Ser	1	\$3,960	\$3,960
CON-SNT-SIP601	SMARTNET 8X5XNBD Multirate 10G IP Ser	1	\$3,960	\$3,960
CON-SNT-SIP601	SMARTNET 8X5XNBD Multirate 10G IP Ser	1	\$3,960	\$3,960
CON-SNT-OC19XFP	SMARTNET 8X5XNBD 1-port OC192/STM64 POS/RPR XFP Optics	1	\$5,000	\$5,000
				\$1,163,240

Table A1-3: Sample Configuration of Egress Router at Data Center