

Personal Web Server

4/4/2003

As we move into the 21st century, we are experiencing the confluence of societal and technological trends which are opening up new opportunities for innovation. In this paper, I describe an opportunity I call “Personal Web Servers.” I’ll begin by describing the societal and technological trends, and the problems with existing solutions. Next, I’ll describe what a Personal Web Server is, and how it can be implemented.

Societal and Technological Trends

For a long time now, the United States has been a nation of nomads. Families move frequently and over long distances, especially for employment opportunities. However, people yearn for “connections” with others – family members, close friends, and people with like interest. Many families and friends share their lives via frequent telephone calls, e-mail, newsletters, and periodic get-togethers, including family reunions. A growing industry is “scrapbooking” where people are taught how to tell their family’s story through the pictures they’ve taken over their lifetime. Many people belong to clubs that focus on their interest, such a woodworking, sewing, collecting, etc.

On the technological front, the innovation which has had the most influence over the past twenty years is the development of the Internet. People have embraced the Internet for personal communications (especially e-mail) and for information. It’s hard to imagine how we got along without the Internet.

Increasingly, families are subscribing to broadband access to the Internet. Dial up access is just too slow and broadband is “always on” allowing you to access the Internet immediately, without the delay of dialing up a modem. And if the family has more than one computer, they most likely have a residential router connected to the broadband modem to share the broadband access and to provide a firewall.

Another major innovation has been digital photography. More and more, people are taking digital pictures and converting their existing paper photos to digital.

Technically advanced people are now creating personal web sites to tell their story. Some are creating family web sites to share their pictures with other family members and friends, especially when they have a new baby. Some create a web site to share a hobby. Others create web logs (blogs) to share their thoughts and comments, and to create a diary of their lives. Some blogs allow all of the extended family members to enter data, creating a single place for sharing personal news.

However, doing this has not been easy. The primary obstacle to setting up a web site is not in creating the web pages – there are many programs available, such as Microsoft’s FrontPage, which make web page creation easy. Templates which provide the complete “look and feel” are available which only require that the web designer add pictures and text. The problem is in hosting the web site – setting up the site on a server.

Non-technical people have no idea what has to be done to set up a web site on a server, and the companies who host web sites are not interested in “small” web sites. The hosting companies would likely have to provide a lot of assistance to the novice web masters, but they would not be able to charge enough to recover their costs within any reasonable time. It might take more than a year to recover just the set up cost for a personal web site. Any additional assistance provided would only add to the cost.

Personal Web Servers

Of course, these hosting problems can be eliminated if someone were willing to offer very low cost web hosting, with essentially unlimited storage, a pre-tailored setup for families (web templates, password protection, etc.), and lots of handholding, but there's no economic incentive for anyone to do that.

The best solution is for the family to set up a server in their home. The cost is a one-time cost for the equipment and they can choose how much processing power and disk space to use. But the complexity of setting up a server on a personal computer is significant, and personal computers are not inexpensive. An additional computer also takes up a lot of physical space.

It is possible to make your home computer a dual use device, functioning as a standard personal computer and as a server but this is even more complex and puts all your computer data at risk. If someone were able to break into your server, they could have access to all your personal files.

What's needed is a physically small, low cost server only device, which is pre-configured to function as a family web server, complete with ID and password security – a device I call the “**Personal Web Server.**”

The Personal Web Server is not a PC – it uses a low cost microprocessor, such as an ARM core, and Flash and RAM memory. It runs a simple operating system and it does only one thing – it functions as a web server. It is EXTREMELY easy to use.

The Personal Web Server can only be effectively used by a residential user who has broadband access (cable modem, DSL, or any other type of broadband access). Additionally, most people who use the Personal Web Server will also have one or more personal computers in their home. While they want the Personal Web Server to be exposed to the Internet, they want to protect their personal computers from access. For this reason, the Personal Web Server should include a router/firewall function as well as the server function. The router/firewall function protects the personal computers in the home, while making the Personal Web Server accessible to others. Products which implement the router/firewall function are available today and are usually called a residential gateway.

Existing residential gateway products contain all of the processing and storage components required to implement the Personal Web server except for storage of the web pages. That is, a web server could be implemented on a residential gateway except that additional storage is required for the web pages. Web pages require a great deal of storage, primarily due to the pictures that people want to include with the web pages – it's fairly easy for a web site to use 50 Mbytes of storage, and more is better. How can this amount of storage be added to a residential gateway without adding cost?

The answer is that it's impossible to add 50 Mbytes or more of storage and not affect the price of the residential gateway. An alternative, however, is to allow the customer to add the storage *after* the sale of the residential gateway. Depending upon the storage selected, the customer may have the storage available already, or it may be very low cost to purchase. But what fits that description?

Removable storage, such as a Compact Flash card, meets these needs. It doesn't have to be Compact Flash, of course. It could be SmartMedia, a MultiMedia card, a Sony memory stick, or any other kind of removable media.



Compact Flash cards are used in digital cameras so many end users may have old Compact Flash cards which are not big enough for their digital camera but could be used for the Personal Web Server. This may especially be true for Compact Flash cards with less than 512 Mbytes.

This means that a residential gateway product can be sold which includes the software for the Personal Web Server, but the web server functionality is not “enabled” until a Compact Flash card is plugged into the socket on the device.

Implementation of the Personal Web Server

The personal web server will consist of the following functional hardware blocks.

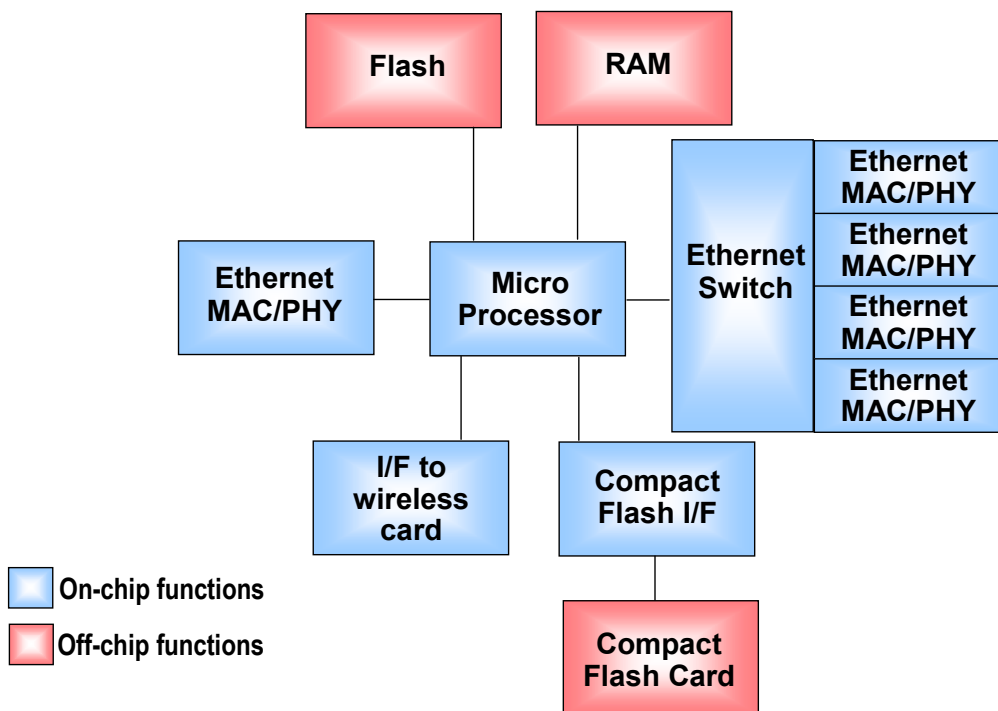


Figure 1: Basic block diagram of the Personal Web Server.

The single Ethernet MAC/PHY on the left side of Figure 1 connects to the broadband interface device (usually a cable or DSL modem). The Ethernet MAC/PHYs on the right side of Figure 1 connect to the personal computers or other devices in the home. The microprocessor implements a router/firewall function which passes packets between the broadband access device and the devices in the home, providing a security function at the same time. A function known as Network Address Translation (NAT) is usually implemented but other techniques could also be used.

The microprocessor also implements the software required for the server functionality, including the ability to respond to a request for web pages, password protection, and file transfer protocol (FTP) which is used to transfer web pages to the server.

Personal Web Server

The software required for the server functionality will add to the flash storage requirements for the product compared to a traditional residential gateway (I estimate that all of the Personal Web Server programming can be contained in about 100K bytes). However, this will not add cost because the existing residential gateways use flash chips which have excessive storage capacity simply because these large capacity flash memories are less expensive than smaller capacity memories. Thus, the primary cost difference between a traditional residential gateway and a Personal Web Server will be interface for the Compact Flash card. This interface consists of logic to create a glueless Compact Flash interface and a socket for the Compact Flash card.

The Compact Flash interface is essentially the same as a PCMCIA interface. Adding the glueless interface to an existing residential gateway chip¹ may not increase the number of pins on the processor chip. For example, the Atmel AT91RM9200 chip implements a glueless interface for CompactFlash, SmartMedia and NAND flash and is shipped in a 208 PQFP package. I have been unable to discover the cost of a Compact Flash socket, but I estimate it at about \$0.50. This means that a residential gateway *with* the Personal Web Server will have a BOM that is about \$0.50 higher than a traditional residential gateway, assuming the residential gateway chip is sold for the same price. Even if this results in a street price that is several dollars higher than a traditional residential gateway, the extra functionality will support the price increase.

Because of this small additional cost, the Personal Web Server can be sold for essentially the same price as a traditional residential gateway. And if the two devices are priced essentially the same, a reasonable purchaser will choose the Personal Web Server even if he/she does not plan to use the server function immediately, because the purchaser may decide to use it in the future.

Note that a Compact Flash card is not included with the Personal Web Server, in order to keep the selling price of the Personal Web Server essentially the same as a traditional residential gateway. Additionally, this allows the purchaser to choose the amount of storage needed and to increase the amount of storage over time to accommodate the increasing number of pictures and text on the web site. For example, a purchaser may have an old Compact Flash card which previously was used in the consumer's digital camera. This card could be used to get started with the server functionality. But even if the consumer has to purchase a new Compact Flash card, the prices continue to drop while capacity increases. Today, a 64MB card is about \$10, a 128 MB is available for about \$20, while a 256 MB card is about \$35 (these prices are all standard prices – not sale prices).

The storage does not have to be Compact Flash, of course. It could be any type of flash card, a hard drive, or even a network mapping of storage on another device (such as a PC). I have used Compact Flash simply because it is common and known to most people.

The software components necessary to implement the Personal Web Server are available from a variety of sources and include an operating system, a TCP/IP stack, a web server function with password security, an FTP server so that pages can be downloaded into the server, a flash file system, a DNS update function (also known as “dynamic DNS client”), and several other functions. One representative supplier of most of these software functions is a company by the name of Interniche (www.interniche.com).

¹ While the lowest cost approach to adding the Compact Flash interface is to incorporate a glueless interface into the residential gateway chip, external logic can be used to provide an interface between the microprocessor bus and the Compact Flash. This may be the best approach to achieve rapid time to market.

Functionality of the Personal Web Server

The Personal Web Server should be designed to be used as a residential gateway and web server, out of the box. The overall functionality should be limited to displaying static web pages – no CGI (Common Gateway Interface) scripts nor forms², for example. Password security should be enabled. A default IP address, probably 192.168.1.1, should be preprogrammed into the unit.

Setup could be accomplished via a password protected web page. In addition to the residential gateway setup, the Personal Web Server setup functions may include the ability to manage a list of IDs and passwords, to enable or disable password access, and to set the IP address for the Personal Web Server function. Statistics on accesses and storage utilization should be available.

The Personal Web Server should have a standard FTP server function so that Web pages can be loaded into the server³.

To begin using the unit, the customer will take it out of the box, provide power, and connect the unit via an Ethernet cable to the cable or DSL modem. The unit would then be setup, probably starting with the residential gateway functions and then with the Personal Web Server functions. This will be done by typing the IP address (192.168.1.1) into the browser on the local PC, and entering the password. The user may wish to change the setup password. The user may also want to set up some IDs and passwords for people to use to access the site. There will be a limit to the number of IDs and passwords, but it can be 50 to 100 sets.

A CD disk can be provided with templates for a family web site. The templates can be used with a web page editor like Microsoft FrontPage, Netscape's editor, or any of the other web page editors. When the first pages are complete, they can be loaded into the server using FTP to the IP address (192.168.1.1) and the setup password. Many web page editors have a "publish" function which is an FTP function built into the program. Alternately, a separate FTP program can be used.

The easiest and cheapest way for the family to give their relatives access to the web site is to send them the broadband IP address (which is available in the residential gateway) and the appropriate ID and password. The relatives simply put the IP address, and perhaps a port number, into their browser and the browser takes care of the rest. The next screen they see will be the security screen asking them to enter the ID and password.

The family is not limited to using the actual IP address – domain names may be used, just like any other web server. In fact, multiple web sites may be hosted on the one IP address if http host headers are supported. This can allow each person in the household to have their own web site (each with unique domain names), all hosted on the Personal Web Server. Multiple sites can also be implemented by using different port numbers for each web site.

Note that if a family gets a domain name, the Personal Web Server could support mail under that domain name. For example, if the family obtained the domain name smithfamily.com, they could have mail addresses of george@smithfamily.com, sally@smithfamily.com, etc. While the Personal Web Server could actually be a POP3 mail host, a simpler approach might be to forward the mail to the mail address

² CGI scripts could be added as an enhancement at some later time, even if the functionality was limited when compared to regular web servers. This would allow people to do a few more things with their servers, such as serving dynamic web pages, handling forms (such as a guest book) and hosting discussion groups.

³ Of course, the Compact Flash card could be removed from the Personal Web Server and plugged into a PC to copy the web pages to the Compact Flash card, then returned to the Personal Web Server – but this is not very convenient.

which is part of the broadband service (like george.smith@cox.net, if Cox was the broadband provider). In this way, the mail addresses on smithfamily.com are essentially aliases to the broadband mail.

What will be on the Personal Web Server?

The primary purpose of the Personal Web Server will be to tell a story. And people have all kinds of stories to tell.

An obvious one is to share pictures of a new baby with friends and family members. A web professional friend of mine recently had a baby and created a web site of pictures of his son. Since he did that, he tells me that he has been inundated with request by friends and relatives to set up “family” web sites for them (he’s resisted the requests so far).

People may want to tell about their hobby. As an example, see www.my427sc.com which is a web site showing the evolution of a hobby car.

People may wish to use the Personal Web Server to host their blog. Blogs are like “instant messages” to the web. See www.blogger.com/about.pyra for an explanation of blogs, and pointers to active blogs. Blogs can be individual or a group of people can all post to the same blog. While most blogs are text, audio blogs are now possible. Pictures can also be included in blogs.

Another application is for a Digital Scrapbook. People will document the lives of their family with digital pictures just as people do today with paper pictures. A paper scrapbook is not just something to look at, but something to share. Two people often look at it together with one telling the other stories about the pictures and the events in the pictures. A digital scrapbook could be viewed in the same way, with communications by phone (or VOIP) and both people viewing the web pages. A special feature of the server could be that one of the people could control the viewing of the pages (the viewing could be “locked” between the people) so that when that person turned to a new web page, that new page would appear on the other person’s browser.

Working parents may want to create a professional web site about themselves, targeted at marketing themselves professionally.

The children may want to create web pages for themselves to share with their school friends, although more immediate communication (phone and IM) may be of greater interest to them.

It’s hard to imagine all the uses web sites might be put to. All that can be said for sure is that people will find many applications as long as the Personal Web Server EXTREMELY easy to use.

Restrictions and Limitations on the Personal Web Server

Although the Personal Web Server is something that many households would like, it is not something that the broadband providers will like or encourage. They view their residential broadband service as being used to access the Internet and not to host web servers. They all provide “business level” broadband services which can be used to connect web servers to the Internet, at a higher cost per month.

This is similar to the attitude broadband provider first took towards residential gateways. They initially prohibited use of their broadband service by more than one computer unless the customer paid a monthly

charge for additional IP addresses. However, customers were simply not willing to pay this when residential gateways were available for a one-time charge. The broadband providers tried several things to limit the use of the residential gateways, including looking for the MAC address of the Ethernet card in the PC, but the residential gateway manufacturers simply provided techniques for spoofing the broadband provider. The broadband providers eventually gave up attempting to prohibit residential gateways and now sell them as part of their installation service. If broadband providers want to put limitations, the best type of limitation is a limit on the amount of data transferred. That's really what the broadband service providers are worried about anyway.

In any case, there are a number of restrictions on the use of the Personal Web Server which may make it unattractive for business use. Due to these restrictions, the Personal Web Server will likely be used ONLY for personal web use⁴, probably accessed by a small group of people. Some of these problems are:

1. Most residential broadband is asymmetric, with more bandwidth to the home than from the home to the Internet. The Personal Web Server will serve pages over the lower speed portion of the broadband link.
2. Some broadband Internet Service Providers block access to port 80, which is used for access to a web server. Several "port scan" web sites are available to check this (see <http://www.securitymetrics.com/portscan.adp>, for example). If port 80 is open, none of the alternate port activities described below need be used.
3. Most broadband customers have a "dynamic" IP address instead of a "static" IP address.

The best situation is to have a static IP address with a broadband service provider who does not block port 80. However, even if a dynamic IP address is provided to the user, and port 80 is blocked, the Personal Web Server can still be used.

The port 80 restriction means that the Personal Web Server must be accessed through a port other than 80. The port used can be essentially any port other than the so-called "well known ports" but a good choice is port 8080⁵. This may require special handling for the family to use a domain name such as www.smithfamily.com.

When someone enters a domain name such as www.smithfamily.com in their browser, the browser makes a request to an Internet service known as domain name services (DNS). The DNS finds the IP address of the http services of the domain (in this example, www.smithfamily.com) and returns the IP address to the browser. From then on, all access to the Smith family server is done through the IP address.

Web access defaults to port 80, which is one of the so-called "well known ports". To access the Smith family server on port 8080, the user would have to enter www.smithfamily.com:8080. The family would have to tell people to use the full name, with the :8080⁶, in order for people to gain access to their web site. If they just used www.smithfamily.com their access would be blocked because of the blocked port 80.

In modern browsers, which implement HTTP⁷ 1.1 or higher, the browser will open a persistent TCP connection when the first contact is made with the server. After this, the browser and server will use the

⁴ A small business which did not have a lot of web hits *could* use the Personal Web Server, especially in the early days of the business. It could later cut over to a standalone server or a web hosting service.

⁵ Generally, any port number above 1024 may be used.

⁶ This will likely not be a problem because family members will bookmark the URL – so they won't have to remember to add the :8080 in later accesses.

⁷ HTTP = HyperText Transfer Protocol

same connection, which will include the proper port (in this example, port 8080). However, there are situations where the TCP connection may be closed. One case is where the person viewing the pages does not request the next web page in a reasonable amount of time. Suppose the person was interrupted while viewing the pages and had to go to a meeting. The web page would be left on the screen but the connection would time out and be closed. When the person returned, he/she would click on a link and expect to be connected to the proper page. The URL, however, would not include the port number and the connection would fail because it would be attempted on port 80.

To problem can be avoided in the Personal Web Server by some special processing of the web page links. To allow web pages to be moved from server to server, links are not coded into the web pages as absolute addresses – they are coded as relative references. That is, the link assumes that the connection has already been established (such as to IP address 68.170.99.22 plus a port number) and the link only references the directories and pages within the computer represented by the connection.

However, since the Personal Web Server knows its own domain name, IP address, and port number that the homeowner wants to use for the server (maybe 8080), it can “process” the web pages and replace the relative references with absolute references that include the domain name (or IP address) and port number⁸. The references to other web pages are easy to identify because they are coded in HTML (HyperText Markup Language), XML (eXtensible Markup Language), JavaScript, or any other standard markup language supported by browsers. Since these are standards, the technique of indicating links is well known for each.

This “processing” can be done in several different ways. One way is to process the web pages as they are loaded onto the server. Another way is to load the pages and then do the “processing” as an offline task. Finally, the pages may be processed when they are served. The best approach may be to process the pages as they are served because if the IP address or port number is changed at any time, the information sent to the browser will be as accurate as possible.

To make access easier, a redirection service can be used. Many companies already provide redirection services, some for free. With this service, the domain name (in this case www.smithfamily.com) would be registered to the redirection server. When a user entered the URL www.smithfamily.com, the user would first access the redirection server which would automatically redirect them to the actual IP address of the Personal Web Server with the port number appended. For example, if the actual IP address of the family’s Personal Web Server was 68.170.99.22, the redirection server would redirect to 68.170.99.22:8080 (or whatever port the family chose to use). This means that when the family’s IP address changes, the redirection server must be notified, as well as the DNS server.

The expense and complexity of setting up a domain name can be avoided if the family only wants to provide access to a small group of people, e.g., their family and close friends. This can be done by providing the family and close friends with the IP address for the home and the port number chosen for the server. For example, suppose my IP address is 68.170.99.22 and I chose to use port 8080 for my server. I would tell my family and friends to type the following into their browser “address line”:
68.170.99.22:8080.

Using this technique is not all that hard for the family and friends because once they make access to the site, they can bookmark the site into their “Favorites” which will remember the address, including the port number. Using this technique avoids the cost of registering a domain name and finding a DNS provider.

⁸ In general, it’s better to use the domain name because someone may bookmark the page. This maintains the link even if the IP address changes.

The second problem is the dynamic IP address. This means that the IP address can change, although in practice it does not change over long periods of time, if at all. Of course, if dynamic DNS is used, the Personal Web Server will simply note that the IP address has changed and report the new IP address to the DNS server. Any subsequent accesses to the domain name will reflect the new IP address. If the redirection service is used, it must also be notified of the IP address change.

If domain names are not being used (if the family just gave the IP address and port number to their friends) the family will have to inform the family and friends of the new IP address.

Both the dynamic IP address problem and redirection can be handled by working with existing companies in the marketplace. Many companies already provide a “Dynamic DNS” service (targeted specifically at this problem of changeable IP address), and most also provide a redirection service. Search on “dynamic dns” and you’ll find a large list. To make these services easier for the residential customer to use, the two services should be combined into one service so that setup is easy for the homeowner.

There are many opportunities for ease of use in this area. DNS is a complex area which the homeowner will certainly not know about and these dynamic DNS companies could offer a “turnkey” solution, handling all of the issues involved in this area, from registering the domain name, to automatically detecting the user’s IP address, to setting up the Dynamic DNS and redirection, all in a single setup session. Of course, the residential gateway manufacturer could provide the redirection and DNS services for an annual fee (maybe free for the first year with purchase of the unit).

Summary

This paper describes the Personal Web Server. The Personal Web Server provides the following advantages to the user:

1. The Personal Web Server can be purchased for approximately the same cost as a traditional residential gateway, meaning that the cost to obtain the additional server functionality is small. The processor power needed to serve static web pages is very small⁹ – the addition of the Personal Web Server function will not impact the residential gateway function.
2. The server is a one-time purchase – no monthly fees beyond that for broadband access need be paid to be able to host a web site.
3. While a domain name can be used, a domain name and the registration and maintenance fees are not required, provided the IP address is distributed for access. This allows the family to get started with the minimum of setup, cost, and complexity.
4. The user can choose the amount of storage for the web pages, and the storage does not have to be purchased until the customer is ready to implement the web server functionality. The user can easily upgrade the amount of storage as necessary. For a home user, 256 megabytes is an enormous amount of web storage, yet it can be purchased for less than \$50 today.

⁹ Web servers handling static web pages have been implemented on an 8MHz 8-bit processor. See <http://www.chipcenter.com/circuitcellar/july99/c79b11.htm> for an article about one. Commercial products on 8-bit processors can be found at www.picoweb.net and www.siteplayer.com.