

# GNU Bayonne: telephony services for freely licensed operating systems

David Sugar <[sugar@gnu.org](mailto:sugar@gnu.org)>  
<http://www.gnu.org/software/bayonne>

## Abstract

GNU Bayonne is a middle-ware telephony server that can be used to create and deploy script driven telephony application services. These services interact with users over the public telephone network. What we are hoping to do is enable, using commodity PC hardware and CTI cards running under GNU/Linux and FreeBSD, which often are available from numerous vendors, to create carrier applications like Voice Mail and calling card systems, as well as enterprise applications such as unified messaging. GNU Bayonne can be used to provide voice response for e-commerce systems and has been used in this role in various e-gov projects. GNU Bayonne can also be used to telephony enable existing scripting languages such as perl and python.

## 1 Introduction

Even without considering all the various reasons of why we must have Free Software as part of the telecommunications infrastructure, it is important to consider what the goals and platform needs are for a telephony platform. Historically, telephony services platforms had been the domain of real-time operating systems. Recent advances in CTI hardware has made it possible to offload much of this requirement to hardware making it practical for even low performance systems running efficient kernels to provide such services for many concurrent users.

Telephony services are usually housed in phone closets or other closed and isolated areas. As such, remote maintainability, and high reliability are both important platform requirements as well. The ability to integrate with and use standard networking protocols is also becoming very important in traditional telephony, and certainly is a key requirement for next generation platforms.

So we can summarize; low latency/high performance kernels, remote manageability without the need for a desktop environment, high reliability, and open networking protocols. This sounds like an ideal match for BSD or Linux kernel based systems.

However, when we looked further into this question and architected GNU Bayonne, we also decided threading was important. Threading offers some interesting design advantages, but, equally important, it provides a means of better utilizing SMP hardware. We found it important to scale up to solutions that can support voice processing on a full DS-3 circuit using a single server. Threading represents some challenges in current BSD systems as I will elaborate further in this paper, so we initially chose to focus primarily on GNU/Linux systems rather than BSD.

Our goal for GNU Bayonne 1.0 was primarily to make telephony services as easy to program and deploy as a web server is today. We choose to make this server easily programmable through server scripting. We also desired to have it highly portable, and allow it to integrate with existing application scripting tools so that one could leverage not just the core server but the entire platform to deliver telephony functionality and integrate with other resources like databases.

GNU Bayonne, as a telephony server, also imposes some very real and unique design constraints. For example, we must provide interactive voice response in real-time. “real-time” in this case may mean what a person might tolerate, or delay of 1/10th of a second, rather than what one might measure in milliseconds in other kinds of real-time applications. However, this still means that the service cannot block, for, after all, you cannot flow control people speaking.

Since each vendor of telephony hardware has chosen to create their own unique and substantial application library interface, we needed GNU Bayonne to sit above these and be able to abstract them. Ultimately we choose to create a driver plug-in architecture to do this. What this means is that you can get a card and API from Aculab, for example, write your application in GNU Bayonne using it, and later choose, say, to use Intel telephony hardware, and still have your application run, unmodified. This has never been done in the industry widely because many of these same telephony hardware manufacturers like to produce their own middle-ware solutions that lock users into their products.

## 2 Supporting Libraries

To create GNU Bayonne we needed a portable foundation written in C++. I wanted to use C++ for several reasons. First, the highly abstract nature of the driver interfaces seemed very natural to use class encapsulation for. Second, I found I personally could write C++ code faster and more bug free than I could write C code.

Why we choose not to use an existing framework is also simple to explain. We knew we needed threading, and socket support, and a few other things. There were no single framework that did all these things except a few that were very large and complex which did far more than we needed. We wanted a small footprint for GNU Bayonne, and the most adaptable framework that we found at the time typically added several megabyte of core image just for the runtime library.

GNU Common C++ (originally APE) was created to provide a very easy to comprehend and portable class abstraction for threads, sockets, semaphores, exceptions, etc. This has since grown into it's own and is now used as a foundation of a number of projects as well as being a part of GNU.

In addition to having portable C++ threading, we needed a scripting engine. This scripting system had to operate in conjunction with a non-blocking state-transition call processing system. It also had to offer immediate call response, and support several hundred to a thousand instances running concurrently in one server image.

Many extension languages assume a separate execution instance (thread or process) for each interpreter instance.

These were unsuitable. Many extension languages assume expression parsing with non-deterministic run time. An expression could invoke recursive functions or entire sub-programs for example. Again, since we wanted not to have a separate execution instance for each interpreter instance, and have each instance respond to the leading edge of an event callback from the telephony driver as it steps through a state machine, none of the existing common solutions like tcl, perl, guile, etc, would immediately work for us. Instead, we created a non-blocking and deterministic scripting engine, GNU ccScript.

GNU ccScript is unique in several ways. It is step executed, and is non-blocking. Statements either execute and return immediately, or they schedule their completion for a later time with the executive. A given "step" is executed, rather than linearly. This allows a single thread to invoke and manage multiple interpreter instances. While GNU Bayonne can support interacting with hundreds of simultaneous telephone callers on high density carrier scale hardware, we do not require hundreds of native "thread" instances running in the server, and we have a very modest CPU load.

Another way GNU ccScript is unique is in support for memory loaded scripts. To avoid delay or blocking while loading scripts, all scripts are loaded and parsed into a virtual machine structure in memory. When we wish to change scripts, a brand new virtual machine instance is created to contain these scripts. Calls currently in progress continue under the old virtual machine and new callers are offered the new virtual machine. When the last old call terminates, the entire old virtual machine is then disposed of. This allows for 100% uptime even while services are modified.

Finally, GNU ccScript allows direct class extension of the script interpreter. This allows one to easily create a derived dialect specific to a given application, or even specific to a given GNU Bayonne driver, simply by deriving it from the core language through standard C++ class extension.

## 3 TGI support and plug-ins

To be able to create useful applications, it is necessary to have more than just a scripting language. It requires a means to be extended so that it can incorporate database access libraries or other functions that fall outside of the scope of the scripting language itself. These extensions should be loaded on demand

only when used, and should be specified at runtime so that new ones can easily be added without the need to recompile the entire server.

To support scripting extensions we have the ability to create direct command extensions to the native GNU Bayonne scripting languages. These command extensions can be processed through plug-in modules which can be loaded at runtime, and offer both scripting language visible interface extensions, and, within the plug-in, the logic necessary to support the operation being represented to the scripting system. These are much more tightly coupled to the internal virtual machine environment and a well written plug-in could make use of thread pools or other resources in a very efficient manner for high port capacity applications.

When writing command extensions, it is necessary to consider the need for non-blocking operations. GNU Bayonne uses ccScript principally to assure non-blocking scripting, and so any plug-in must be written so that if it must block, it does so by scheduling a state operation such as "sleep" and performs potentially blocking operations in separate threads. This makes it both hard and complex to correctly create script extensions in this manner.

While GNU Bayonne's server scripting can support the creation of complete telephony applications, it was not designed to be a general purpose programming language or to integrate with external libraries the way traditional languages do. The requirement for non-blocking requires any module extensions created for GNU Bayonne are written highly custom. We wanted a more general purpose way to create script extensions that could interact with databases or other system resources, and we choose a model essentially similar to how a web server does this.

The TGI model for GNU Bayonne is very similar to how CGI works for a web server. In TGI, a separate process is started, and it is passed information on the phone caller through environment variables. Environment variables are used rather than command line arguments to prevent snooping of transactions that might include things like credit card information and which might be visible to a simple "ps" command.

The TGI process is tethered to GNU Bayonne through stdout and any output the TGI application generates is used to invoke server commands. These commands can do things like set return values, such as the result of a database lookup, or they

## Bayonne Architecture

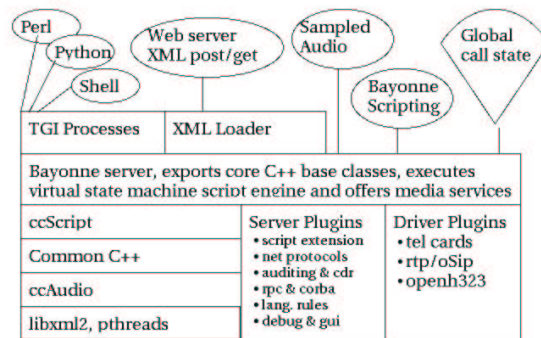


Figure 1: Architecture of GNU Bayonne

can do things like invoke new sessions to perform outbound dialing. A "pool" of available processes are maintained for TGI gateways so that it can be treated as a restricted resource, rather than creating a gateway for each concurrent call session. It is assumed gateway execution time represents a small percentage of total call time, so it is efficient to maintain a small process pool always available for quick TGI startup and desirable to prevent stampeding if say all the callers hit a TGI at the exact same moment.

## 4 Bayonne Architecture

As can be seen, we bring all these elements together into a GNU Bayonne server, which then executes as a single core image. The server itself exports a series of base classes which are then derived in plug-ins. In this way, the core server itself acts as a "library" as well as a system image. One advantage of this scheme is that, unlike a true library, the loaded modules and core server do not need to be relocatable, since only one instance is instantiated in a specific form that is not shared over arbitrary processes.

When the server comes up, it creates gateways and loads plug-ins. The plug-ins themselves use base classes found in the server and derived objects that are defined for static storage. This means when the plug-in object is mapped through dload, it's constructor is immediately executed, and the object's base class found in the server image registers the object with the rest of GNU Bayonne. Using this method, plug-ins in effect automatically register themselves through the server as they are loaded, rather than through a separate

runtime operation.

The server itself also instantiates some objects at startup even before main() runs. These are typically objects related to plug-in registration or parsing of the configuration file.

Since GNU Bayonne has to interact with telephone users over the public telephone network or private branch exchange, there must be hardware used to interconnect GNU Bayonne to the telephone network. There are many vendors that supply this kind of hardware and often as PC add-on cards. Some of these cards are single line telephony devices such as the Quicknet LineJack card, and others might support multiple T1 spans. Some of these cards have extensive on-board DSP resources and TDM busses to allow interconnection and switching.

GNU Bayonne tries to abstract the hardware as much as possible and supports a very broad range of hardware already. GNU Bayonne offers support for /dev/phone Linux kernel telephony cards such as the Quicknet LineJack, for multiport analog DSP cards from VoiceTronix and Dialogic, and digital telephony cards including CAPI 2.0 (CAPI4Linux) compliant cards, and digital span cards from Intel/Dialogic and Aculab. We are always looking to broaden this range of card support.

At present both voice modem and OpenH323 support is being worked on. Voice modem support will allow one to use generic low cost voice modems as a GNU Bayonne telephony resource. The openh323 driver will actually require no hardware but will enable GNU Bayonne to be used as an application server for telephone networks and softswitch equipment built around the h323 protocol family. At the time of this writing openh323 support is slated for release as part of GNU Bayonne 1.1.

## 5 GNU Bayonne and XML Scripting

Some people have chosen to create telephony services through web scripting, which is an admirable ambition. To do this, several XML dialects have been created, but the idea is essentially the same. A query is made, typically to a web server, which then does some local processing and spits back a well formed XML document, which can then be used as a script to interact with the telephone user. These make use of XML to generate application logic and control much like a

scripting language, and, perhaps, is an inappropriate use of XML, which really is designed for document presentation and inter-exchange rather than as a scripting tool. However, given the popularity of creating services in this manner, we do support them in GNU Bayonne.

GNU Bayonne did not choose to be designed with a single or specific XML dialect in mind, and as such it uses a plug-in. The design is implemented by dynamically transcoding an XML document that has been fetched into the internal ccScript virtual machine instructions, and then execute the transcoded script as if it were a native ccScript application. This allows us to transcode different XML dialects and run them on GNU Bayonne, or even support multiple dialects at once.

Since we now learn that several companies are trying to force through XML voice browsing standards which they have patent claims in, it seems fortunate that we neither depend on XML scripting nor are restricted to a specific dialect at this time. My main concern is if the W3C will standardize voice browsing itself only to later find out that the very process of presenting a document in XML encoded scripting to a telephone user may turn out to have a submarine patent, rather than just the specific attempts to patent parts of the existing W3C voice browsing standard efforts.

Currently GNU Bayonne implements a "BayonneXML" dialect as a model XML plugin. This dialect demonstrates a range of functionality similar to "CallXML". We have had offers from various sources to fund specific development of W3C spec compliant XML dialects, but so far none of these offers have ever reached the point where a check was cut. It would take considerable time and talent to finish GNU Bayonne XML work, and none of the people actively using it have pushed for XML support. As such, it has received a lower profile in the list of features we wish to currently work on.

## 6 Current Status

The 1.0 release of GNU Bayonne was released on September 1st. This release represents several years of active development and has been standardized in how it operates and how it is deployed. This release is part of the GNU project and has been packaged for use with many GNU/Linux distributions. With this release we have had a stable platform for developing

GNU Bayonne applications and for considering future development.

GNU Bayonne does not exist alone but is part of a larger meta-project, "GNUCOMM". The goals of GNUCOMM is to provide telephony services for both current and next generation telephone networks using freely licensed software. These services could be defined as services that interact with desktop users such as address books that can dial phones and softphone applications, services for telephone switching such as the IPSwitch GNU softswitch project and GNU oSIP proxy registrar, services for gateways between current and next generation telephone networks such as troll and proxies between firewalled telephone networks such as Ogre, realtime database transaction systems like preViking Infotel and BayonneDB, and voice application services such as those delivered through GNU Bayonne.

## 7 GNU Bayonne and FreeBSD

GNU Bayonne is successfully used with GNU/Linux systems today. It is widely used in many areas ranging from commercial carriers in Europe to state governments in the United States. We do not believe telephony should be restricted to any one platform, however, and, even from the beginning, choose to make GNU Bayonne highly portable.

The core libraries that compose GNU Bayonne are all highly portable, and in particular, are built on a single abstract interface library, GNU Common C++. GNU Common C++ offers portable threading and socket support, and has been ported to many platforms, including FreeBSD, as well as some non-Unix platforms. These libraries also have active support and are distributed with build files for directly building BSD style "ports" collection entities. As such it is very simple to build a "port" of Common C++, ccAudio, or ccScript. One just does "make ports" from the master Makefile after using ./configure.

I do not actively package GNU Common C++, ccAudio, or ccScript for FreeBSD just as I do not actively provide Debian package for the GNU/Debian distribution, or for any other target OS. At one time I did many of these things, but I found it became too overwhelming to both manage and build releases, and to personally provide binary and build packages for every target platform. We also do not have a large enough group of active developers where one can simply

work on packaging. In fact, we now depend on the broader GNU/Linux community and vendors for packaging of GNU Bayonne for specific GNU/Linux distributions and for patches when needed for specific distributions.

I actually do on occasion develop under FreeBSD. I actually authored the original FreeBSD port builds for GNU Bayonne's supporting libraries and tested them on my FreeBSD development system at home. I also did an experimental build of GNU Bayonne under FreeBSD and we are hoping to demonstrate it at EuroBSD. However, we have ran into several issues in building GNU Bayonne under FreeBSD related to threading.

The primary issue with FreeBSD threading boils down to two issues; one, that not all blocking system calls behave as cancellation points, and two, that "immediate" cancellation is not at all supported in FreeBSD's native libc\_r threading environment. This causes a number of abhorrent behaviors on the test server I have built under FreeBSD (4.6), which at this time almost works correctly.

There is also the port of the "LinuxThreads" package available for FreeBSD. I recently modified GNU Common C++ to support this as an optional build choice under FreeBSD. The "LinuxThreads" package, I gather, uses the FreeBSD version of "clone()", and impliments posix threads the way that Linux glibc does. The question I have, and I hope to resolve by talking with active FreeBSD developers, is if it makes sense to use and require "LinuxThreads" for supporting GNU Bayonne in the future, or if it makes more sense to work out the issues that prevent the native threading library from working.

The final challenge we have is that there is very limited computer telephony hardware choices available for xBSD systems today.

## 8 EuroBSD conference goals

One goal we have is to generate more interest in the BSD community in general about telephony and in particular about GNU Bayonne. We are looking for help from the BSD community in several areas. One area is to find a person to help coordinate distribution and updates of the FreeBSD "ports" collections for GNU Bayonne and dependent packages.

We also wish to discuss the various issues related to threading and various BSD platforms. In that the current GNU Bayonne developers have fairly limited exposure to xBSD events, there is a lack of full understanding of these issues in xBSD and how they relate to the upcoming 5.0 release of FreeBSD.

Ideally we would like to have contributions from the BSD community as part of GNU Bayonne. In part, such contributions would be particularly helpful in making GNU Bayonne better able to build and be used on the various BSD systems. However, contributions can take many forms, not all of which are coding.

Finally, we wish to further interest computer telephony oem's to support BSD platforms, There are many performance advantages to the BSD kernel, and many reasons it would be useful for such vendors to target development under BSD as well as under the Linux kernel. Some vendors, such as Voicetronix, already actively do this with freely licensed drivers that work on FreeBSD as well as Linux kernels. Most cti vendors neither have freely licensed drivers nor choose to support BSD at all.

## 9 Future Development

While XML is not an immediate area of active development, we are working in several important areas for the next major release. These include designing support for building a complete script driven PBX/telephone switch with GNU Bayonne. Such a system would offer direct application integration with a complete GNU Bayonne hosted phone system that can be used by a small office. This work is initially possible with the new line of Voicetronix PBX cards and their freely licensed drivers for both GNU/Linux and FreeBSD systems.

We are also looking to expand support for additional telephony hardware. In particular, we are interested in completing support for the Zapata telephony interfaces for the 1.1 release. These drivers are available for both GNU/Linux and FreeBSD systems and are freely licensed. The Zapata line includes support for digital (T1) voice resource cards as well as analog telephony hardware. I see this as the first opportunity for FreeBSD systems to participate in high density digital telephony solutions such as those we use GNU/Linux for with GNU Bayonne today with commercial carriers.

Certainly, one of our most important goals is making GNU Bayonne fully available for BSD systems, and I hope that will finally be accomplished as part of the upcoming 1.1 release.

## 10 Acknowledgments

There are a number of contributors to GNU Bayonne. These include Matthias Ivers who has provided a lot of good bug fixes and new scheduler code. Matt Benjamin has provided a new and improved TGI tokenizer and worked on Pika out-bound dialing code. Wilane Ousmane helped with the French phrasebook rule sets and French language audio prompts. Henry Molina helped with the Spanish phrasebook rule sets and Spanish language audio prompts. Kai Germanschewski wrote the CAPI 2.0 driver for GNU Bayonne, and David Kerry contributed the entire Aculab driver tree. Mark Lipscombe worked extensively on the Dialogic driver tree. There have been many additional people who have contributed to and participated in related projects like GNU Common C++ or who have helped in other ways.