Reprinted from the

Proceedings of the Ottawa Linux Symposium

June 26th–29th, 2002
Ottawa, Ontario
Canada

Conference Organizers

Proceedings Formatting Team

GNU Bayonne: telephony application server of the GNU project

David Sugar
Open Source Telecom.
Somerset, NJ, 08873
sugar@gnu.org
http://www.gnu.org/software/bayonne

Abstract

GNU Bayonne is a middleware 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. GNU Bayonne can be used to create carrier applications like Voice Mail and calling card systems, as well as enterprise applications such as unified messaging. It 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

Our goal in GNU Bayonne was to make telephony services as easy to program and deploy as a web server is today. We choose to make this server easily programmable thru 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. “realtime” 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 plugin 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 middleware solutions that lock users into their products.

2 GNU Common C++

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 Bayonne, and the most adaptable framework that we found at the time typically added several megs 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 its own and is now used as a foundation of a number of projects as well as being a part of GNU.

3 GNU ccScript

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 subprograms 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 thru 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 vm and new callers are offered the new vm. When the last old call terminates, the entire old vm 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 thru standard C++ class extension.

4 TGI support and plugins

While GNU Bayonne offers a ccScript virtual interpreter for creating telephony applications, we wanted to be able to integrate support for databases and other things. There are systems and scripting environments such as Perl which already offer database connectivity. So we created a concept called "TGI," which, like CGI, allows external executables to be invoked from within a call flow script, and the results to be recorded so that information can be passed both to and from the user.

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 thru 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 thru stdout and any output it 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 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.

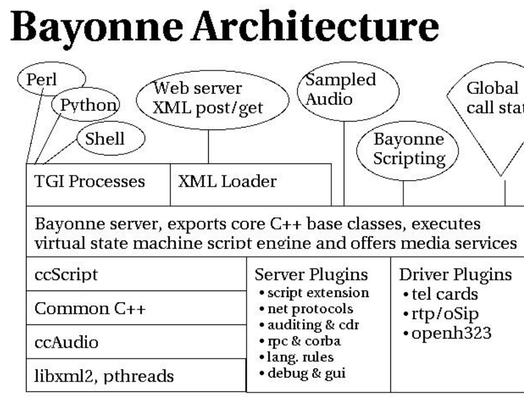

Figure 1: Architecture of GNU Bayonne

TGI does involve a lot of overhead, and so in addition we have the ability to create direct command extensions to the native GNU Bayonne scripting languages. These command extensions can be processed thru plugin modules which can be loaded at runtime, and offer both scripting language visible interface extensions, and, within the plugin, 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 plugin could make use of thread pools or other resources in a very efficient manner for high port capacity applications.

5 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 plugins. 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 plugins. The plugins themselves use base classes found in the server and derived objects that are defined for static storage. This means when the plugin object is mapped thru dload, its 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, plugins in effect automatically register themselves thru the server as they are loaded, rather than thru a separate runtime operation.

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

6 Hardware Requirements

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 I am not sure if either or both of these will be completed in time for the 1.0 release.

7 GNU Bayonne and XML Scripting

Some people have chosen to create telephony services thru web scripting, which is an admerable 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 plugin. The design is implimented 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 thru 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.

8 Current Status

At the time of this paper's publication, the 1.0 release of GNU Bayonne should already be in active distribution. This release represents several years of active development and has been standardized in how it operates and how it is deployed. Even before this point, and for the past 6 months, active development has happened on a second generation GNU Bayonne server, and snapshots of this new server are currently available for download. Where GNU Bayonne is evolving will be explained further on.

9 GNU Bayonne the Meta Projects

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 thru GNU Bayonne.

10 Transactional Databases

BayonneDB is mentioned briefly for transactional services. When we conceived of the need for a transactional database server, we considered that database queries might be slow. The telephony server does not want to do nothing while a transaction is completing, especially if it takes many seconds to happen. Maybe the caller needs to be played music on hold or given other options.

To accomplish non-blocking transactions that

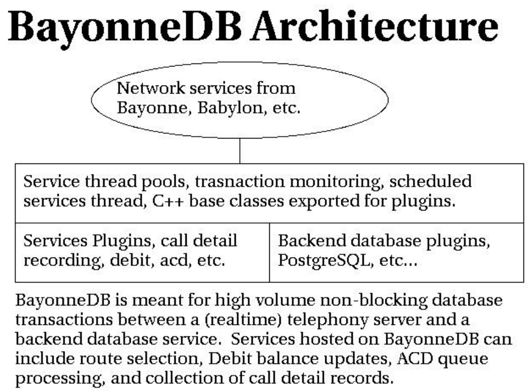

Figure 2: Architecture of BayonneDB

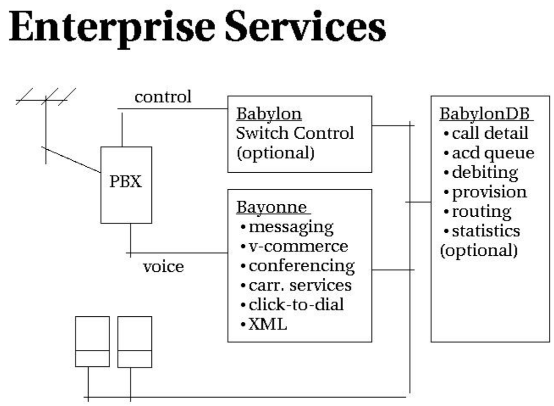

Figure 3: Enterprise Applications Today

allow the telephony server to continue call processing, we choose a peer messaging architecture. A request would be sent to an external server for a transaction, and when the transaction completes, a result message would be sent to the server. There can be time-out and retransmission controls which allow this to be conduced thru UDP packets rather than potentially blocking TCP sessions. This set of protocols and specifications was created initially by Zaheer Milari and myself and published early last year.

BayonneDB was an attempt to implement the concepts in an operational server. Like Bayonne, BayonneDB offers abstraction thru plugins and is based on GNU Common C++. In the case of BayonneDB, it is designed to abstract the interface to the underlying database server used to complete the transaction request. Being threaded, BayonneDB can maintain a persistent threadpool of database connections to optimize overall query performance. A short diagram of BayonneDB architecture is presented below:

11 Enterprise Applications

In our broadest view of enterprise telephony applications, we can see using GNU Bayonne as a part of an overall solution. GNU Bayonne must be able to interact with enterprise data, whether thru transaction monitors such as BayonneDB or thru perl scripts executed via TGI. It may need to interact with other services such as email when delivering voice messages to a unified mailbox, or the local phone switch thru a resource such as Babylon. We will explain Babylon a bit later.

Our view of GNU Bayonne and telephony application services are that it is a strategic and integral part of the commercial enterprise. Proprietary solutions that are in common use today have often been designed from the question of how to lock a user into a specific OEM product family and control what a user or reseller can do or integrate such products, rather than from the question of what the enterprise user needs and how to provide the means to enable it. This has often kept telephony separate and walled off from the rest of the enterprise. We do not wish to see it separate but a natural extension, whether of web services, of contact manage-

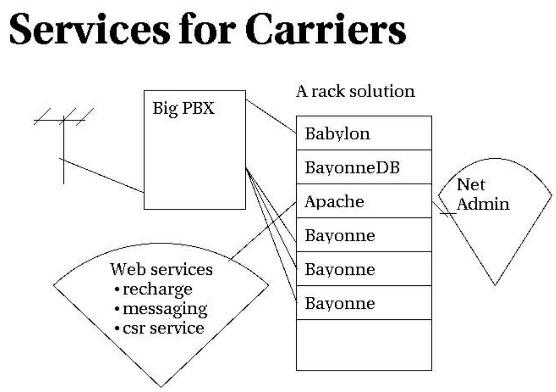

Figure 4: Carrier Applications Today

ment, of customer relations, etc.

12 Carrier Applications

When we look at carrier class applications for GNU Bayonne today, we typically consider applications like operator assistance systems, prepaid calling platforms, and service provider voice mail. Each of these has different requirements. What they have in common is that a front end central office switch might be used, such as a Lucent Excel or even a full ESS 5 switch. Application logic and control for voice processing would then be hosted on one or more racks of GNU Bayonne servers most likely interconnected with multiple T1 spans. If database transactions are involved, such as in pre-paid calling, perhaps we would distribute a BayonneDB server to provide database connectivity for each rack. A web server may also exist if there is some web service component.

Operator assist services are probably the easiest to understand. Very often a carrier might need to provide directory assistance or some other form of specialized assist service. A call will come in from the switching center to a GNU Bayonne server, which will then decide what to do with the call. If the caller is from a location that is known, perhaps the call will be re-routed by GNU Bayonne thru an outgoing span to a local service center. Online operator assistance might be done by creating an outgoing session to locate an operator and then bridge the callers, all on a GNU Bayonne server.

In service provider voice mail one doesn't have to bridge calls. Service provider voice mail is typically much simpler than enterprise voice mail; there is no company voice directory, there is no forwarding or replying between voice mailboxes, there may be no external message notification. All these things make it an easy to define application on first apperance. What it must be is reliable, and ideally scalable.

The problem with service provider voice mail is where to store the potentially large pool of message boxes. We don't know what callers might call in for messages or when. If the call capacity is larger than a single server can handle even with multiple T1 spans, then we might need to deal with a reliable message store hosted on a machine outside the GNU Bayonne servers. We could also scatter mailboxes over multiple machines by hashing the mailbox address into a GNU Bayonne server address, and load balance over multiple servers that way.

If we have a common external message store, perhaps we can have it on a fibre channel bus. GNU Bayonne doesn't like blocking, and traditional network file systems, like NFS, can have long timeout and blocking intervals. Messages can also be transported from a central store over different procotols. One thought I had was a UDP based transport protocol for voice messaging. Since the need is not for full duplex voice, many of the issues in regard to latency and packet size can be relaxed for transporting a voice stream over what is typically required to make VoIP systems work. With a network addressable message store, GNU Bayonne can provide a reliable platform for service provider voice mail.

Many applications carriers wish to deploy do not nessisarly require “carrier grade” Linux to appear before they can be used. In fact, IDT Corp, a major provider of prepaid calling in the world today, uses over 500 rack mounted commodity PC’s running things including a standard distribution of “RedHat” GNU/Linux to reliably service over 20 million call minutes per day in their main switching center. This does not mean there is no value in the carrier grade kernel work, just that it is not nessisary to create and sell some types of GNU/Linux voice processing solutions for carriers today. We have looked at the issues involved in high reliability/carrier grade enhanced Linux and we intend to address those as described a little further.

13 GNU Bayonne clustering

In England one enterprising fellow is working on GNU Bayonne tandem switching nodes. A tandem switching node essentially routes call traffic between spans based on various rules, perhaps to achieve a low interconnection count or to find the least cost available route in a telephone network. This touches upon an interesting and unique feature of GNU Bayonne which we have not yet talked about; GNU Bayonne servers talk to each other.

When Bayonne servers talk with each other, they do two things. Each node elects a “buddy” node to act as a failover for itself. Elections are held every few minutes and the design of this is that a single node will only elect itself to buddy up to two additional nodes in the network. Buddies are useful in failover, since they are aware of all transactions and the state of each GNU Bayonne server, and can complete transactions if a given machine (node) goes down. By having a limited set of buddies chosen thru election, we assure there is no network stampede when a node goes down on the part of other nodes wishing to complete transactions for it.

Since global call state is shared among GNU Bayonne servers, each server knowns what the other one is doing and what its current utilization is like. This can be very useful in a tandem switching application where one needs to know where available endpoints are and if there are ports available at each end point for a given call request. GNU Bayonne cluster networking is still in its infancy, and we are looking for ways to express networking thru the application scripting language.

The main use of clustering at the moment is to overcome the inherit limits of system reliability for acceptance of GNU Bayonne in developing carrier class applications. Over time, this need will be lessened as we take advantage of the work being done on carrier grade GNU/Linux.

14 The NG Server

Even before GNU Bayonne 1.0 had been finalized, work had been started by late last year on a successor to GNU Bayonne. This successor attempts to simplify many of the architectural choices that were made early on in the project to make it easier to adept and integrate GNU Bayonne in new ways.

One of the biggest challegnes in the current

GNU Bayonne server is the creation of telephony card plugins. These often involve the implementation of a complete state machine for each and every driver, and very often the code is duplicated. GNU Bayonne “2” solves this by pushing the state machine into the core server and making it fully abstract thru C++ class extension. This allows drivers to be simplified, but also enabled us to build multiple servers from a single code base.

Another key difference in GNU Bayonne “2” is much more direct support for carrier grade Linux solutions. In particular, unlike GNU Bayonne, this new server can take ports in and out of service on a live server, and this allows for cards to be hotplugged or hot swapped. In a carrier grade platform, the kernel will provide notification of changeover events and application services can listen for and respond to these events. GNU Bayonne is designed to support this concept of notification for management of resources it is controlling.

Finally, GNU Bayonne “2” is designed from the ground up to take advantage of XML in various ways. It uses a custom XML dialect for a configuration language. It also acts as a web service with both the ability to request XML content that describe the running state of GNU Bayonne services and the ability to support XMLRPC. This fits into our vision for making telephony servers integrate with web services, and will be described further in a seperate paper.

15 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 tokeniser and worked on Pika outbound dialing code. Wilane Ousmane helped with the French phrasebook rulesets and French language audio prompts. Henry Molina helped with the Spanish phrasebook rulesets 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 extensivily 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.