



PUMA

Reliable, secure, scalable and user-oriented design of a multi platform framework based on the most advanced stage of web technologies *

Giacomo Strangolino Lucio Zambon

Elettra Sincrotrone Trieste

* inspired by an idea of Alessio I. Bogani



HISTORY

from canone to canone3 (PUMA)

- 2006 Canone Python server, drag and drop designer
- 2016 ElettrApp Responsive, Cordova, JQuery, Bootstrap
- 2017 PWMA C++ server + Websocket (named Canone 2), React, React Native
- <u>2020</u> PUMA NChan server + SSE (named Canone 3), interactive SVG, designer tool for responsive pages
- <u>2022</u> PUMA NChan + SSE (caserver) more robust, modular and distributed architecture





SECTION I

THE SERVICE





BUREAU VERITAS

DESIGN RATIONALE

1. RELIABLE

Websocket issues with proxies...

Events, stream

over channels

The system must work

- From any place, time, platform
- Always, regardless the number of clients
- Always, included when part of the system is unavailable
- Included when network performance is slow (or subject to charges!)





2. SECURE

- 1) From external attacks (DoS, intrusion)
- 2) From ill designed clients flooding the service
- 3) Protect the downstream Control System Engine (Tango, EPICS)

Network architecture
Security-oriented **OS**es (BSD!?)
Framework design, database and code

No service performance decay
 Hamper unruly clients



SSE! :-)



2. SECURE (II)

<u>GOAL</u> secure enough as to avoid VPN or other additional client-side configuration hindering *usability*, especially on mobile devices







3. SCALABILITY

A good infrastructure must be designed with scalability in mind. It reinforces security and reliability.

- 1) <u>Horizontal</u>
- 2) Vertical







3. SCALABILITY (II)







4. GENERIC API

Must serve

- 1) The web
- 2) Mobile applications
- 3) Desktop applications *

* Cumbia libs already support the API so that any Qt application can be instructed to rely on either the native control system engine or the HTTP service at runtime





Nginx + nchan + http/SSE * = ?

* inspired by an intuition of Alessio I. Bogani





Nginx + nchan + http/SSE = ?

- ✓ NGINX: high performance load balancer, web server and reverse proxy *https://www.nginx.com/*
- ✓ **NCHAN**: flexible *pubsub* for the modern web
- ✓ SSE: a server *push* technology: a client receives automatic updates from a server via HTTP connection

A scalable, secure, efficient service with multiplexing for web, mobile and desktop applications





Nginx + Nchan + Redis

Redis (Remote Dictionary Server) is an in-memory data structure project implementing a distributed, in-memory key–value database with optional durability.

✓ Redis can be used to add data persistence and horizontal scalability, failover and high availability to a Nchan setup.

✓ Redis Cluster provides a way to run a Redis installation where data is automatically sharded across multiple Redis nodes.

Nchan + Redis

- \checkmark add scalability via sharding channels among cluster nodes.
- ✓ Redis cluster provides automatic failover, high availability,
- ✓ Redis cluster eliminates the single point of failure of one shared Redis server





Nginx + nchan + http/SSE = ?



- Additionally
- Synchronous readings
- Synch database property fetch
- Authenticated synchronous writings

Multiplexing: *n* clients reading $\underline{x} \longrightarrow 1$ reader to the native engine



Nginx + nchan – scalability and load balancing



Elettra

Sincrotrone Trieste

Nchan + Redis = additional scalability



NEW SERVICE ARCHITECTURE – v1.3.6 (I)

Monitoring over time is a complementary task compared to a single (first) read and configuration phase.

- Monitoring is asynchronous by nature (channels)
- Set up is synchronous: a client subscribes using an *http* request and expects some kind of reply, possibly with *source* configuration and current value.





DESIGN RATIONALE – NEW SERVICE ARCHITECTURE – v1.3.6 (II)

Splitting a single one into two services mirrors real functionality:

- A single service is dedicated to a specific task
- Each single service is *much simpler* because it does not need implement *complementary* operations
- Each service is *much* easier to maintain, and so is debugging
- caserver-async and caserver-sync are the respective names
- A third, intermediate service is needed to proxy (split) the client requests and forward them to either or both services (caserverproxy)









Caserver new architecture (I)

Subscribe requests from clients (method: s) imply

 \checkmark an immediate synchronous reply containing the value and configuration of the source

 \checkmark a subscription for updates through an nchan channel on which the client shall listen

 \checkmark a subscribe request shall be forwarded both to the sync and async caserver instances.

✓ The replies from the caservers are then merged together by caserver proxy and sent back to the client as an http reply to the client's http request.





Caserver new architecture (II)

Handle unsubscribe requests (method: u)

 ca-proxy forwards requests to <u>all</u> known caservers because it's unaware of which one is monitoring the source (due to nginx load balancing across multiple upstream services)

Handle ungracious / vanished clients

 clients must periodically send a keepalive message in order to receive data from channels;

ca-proxy shall periodically check a database to stop monitoring sources linked to vanished clients.





Read requests are forwarded to caserver-sync only





DESIGN RATIONALE – NEW SERVICE ARCHITECTURE – v1.3.6 (IV)

subscribe only requests



subscribe-only requests are forwarded solely to <u>caserver-async</u>





DESIGN RATIONALE – NEW SERVICE ARCHITECTURE – v1.3.6 (V)

unsubscribe requests



monitoring src: <u>broadcast</u> unsubscribe





ca-proxy checks for missing clients (not updating their keep-alive messages)





DESIGN RATIONALE - DEPLOYMENT







TEST ENVIRONMENT – LOAD BALANCING

The test environment is pictured in the previous slide. A query to the database shows the the distribution of 1255 readings across six services:

id	started	expected	state addr	count
311	2021-09-13 13:35:18.331211	2021-09-13 14:20:23.492731	+ ACTIVE 192.168.205.159	212
320	2021-09-13 13:35:17.06742	2021-09-13 14:20:30.02976	ACTIVE 192.168.205.106	222
326	2021-09-13 13:35:17.653089	2021-09-13 14:20:21.798637	ACTIVE 192.168.205.149	235
327	2021-09-13 13:35:17.54277	2021-09-13 14:20:30.455891	ACTIVE 192.168.205.158	215
324	2021-09-13 13:35:16.584578	2021-09-13 14:20:20.77197	ACTIVE 192.168.205.157	161
328	2021-09-13 13:35:18.383527	2021-09-13 14:20:26.159456	ACTIVE 192.168.205.166	210
(6 row	vs)			

1255





TEST ENVIRONMENT – LOAD BALANCING (II)

Load balancing is accomplished by *nginx*. Take a look at *nginx.conf*:

```
http {
    upstream caserver {
        hash $http_x_channel consistent;
        A request is sent to the server with the least number of active connections
        https://docs.nginx.com/nginx/admin-guide/load-balancer/http-load-balancer/#c
        least_conn;
        server puma-01.elettra.trieste.it:9292 fail_timeout=20s;
        server puma-02.elettra.trieste.it:9292 fail_timeout=20s;
        server puma-03.elettra.trieste.it:9292 fail_timeout=20s;
        server puma-04.elettra.trieste.it:9292 fail_timeout=20s;
        server puma-05.elettra.trieste.it:9292 fail_timeout=20s;
        server puma-05.elettra.trieste.it:9292 fail_timeout=20s;
        server puma-dev.elettra.trieste.it:9292 fail_timeout=20s;
    }
}
```





TEST ENVIRONMENT – FAILOVER

Simulate a server failure in one of the hosts:

puma@puma-04:~\$ host puma-04 puma-04.elettra.eu has address 192.168.205.159 puma@puma-04:~\$ sudo killall -SEGV caserver

Take a look at the casupervisor logs:

[Mon	Sep	13 1	4:44:17	2021]	[thread:0x7f7fc717d700]	ca-supervisor:	#	ACTIVE	"pwma-dev" 192.168.205.106:9292	expected:	13s	[Mon Sep 13	15:44:3	0 2021]	started:	Sun Sep	13 13:35:	17 2021
- [Mon	Sep	13 1	4:44:17	2021]	[thread:0x7f7fc717d700]	ca-supervisor:	#	ACTIVE	"puma-01" 192.168.205.149:9292	expected:	15s	[Mon Sep 13	15:44:3	2 2021]	started:	Sun Sep	13 13:35:	17 2021
[Mon	Sep	13 1	4:44:17	2021]	[thread:0x7f7fc717d700]	ca-supervisor:	#	ACTIVE	"puma-02" 192.168.205.157:9292	expected:	14s	[Mon Sep 13	15:44:3	1 2021]	started:	Sun Sep	13 13:35:	16 2021
[Mon	Sep	13 1	4:44:17	2021]	[thread:0x7f7fc717d700]	ca-supervisor:	#	ACTIVE	"puma-03" 192.168.205.158:9292	expected:	14s	[Mon Sep 13	15:44:3	1 2021]	started:	Sun Sep	13 13:35:	17 2021
[Mon	Sep	13 1	4:44:17	2021]	[thread:0x7f7fc717d700]	ca-supervisor:	#	ACTIVE	"puma-05" 192.168.205.166:9292	expected:	19s	[Mon Sep 13	15:44:3	6 2021]	started:	Sun Sep	13 13:35:	18 2021
[Mon	Sep	13 1	4:44:17	2021]	[thread:0x7f7fc717d700]	ca-supervisor:		ZOMBIE	"puma-04" 192.168.205.159:9292	expected:		[Mon Sep 13	15:42:3	4 2021]	started:	Sun Sep	13 13:35:	18 2021





TEST ENVIRONMENT – FAILOVER (II)

Execute the same query as in the first *load balancing* inspection:

id	started	expected state addr	count
320 2021-0	9-13 13:35:17.06742	++++++++	05.106 434
326 2021-0	9-13 13:35:17.653089		05.149 235
327 2021-0	9-13 13:35:17.54277		05.158 215
324 2021-0	9-13 13:35:16.584578	2021-09-13 14:47:01.49247 ACTIVE 192.168.20)5.167 161
328 2021-0	9-13 13:35:18.383527	2021-09-13 14:46:56.885659 ACTIVE 192.168.20)5.166 210

1255

Please note that load redistribution after a failure is

always administered by nginx





TEST ENVIRONMENT

Load balancing and service status: web view



service history activity history

id	started	expected	state	conf_id	srvnam	addr	port
320	2021-09-15 10:44:12.239299	2021-09-15 10:45:08.870399	ACTIVE	1	pwma-dev	192.168.205.106	9292
317	2021-07-20 10:09:22.534603	2021-07-20 12:27:27.964855	INACTIVE	4	pwma-dev	192.168.205.106	19292
326	2021-09-15 10:44:13.767563	2021-09-15 10:45:05.899086	ACTIVE	5	puma-01	192.168.205.149	9292
324	2021-09-15 10:44:12.994295	2021-09-15 10:45:11.993982	ACTIVE	7	puma-02	192.168.205.157	9292
333	2021-09-14 16:23:37.325481	2021-09-14 16:24:48.04815	ACTIVE	9	puma-03	192.168.205.158	9292
332	2021-09-14 16:23:37.907864	2021-09-15 10:45:11.991821	ACTIVE	10	puma-04	192.168.205.159	9292
328	2021-09-14 16:23:38.346319	2021-09-15 10:45:10.384647	ACTIVE	12	puma-05	192.168.205.166	9292
300	2021-07-06 14:53:22.219985	2021-07-06 15:05:30.619484	INACTIVE	18	woody	192.168.1.159	9292





DESIGN RATIONALE (II)

SECTION II

NATIVE CLIENTS





DESIGN RATIONALE (II)

1. MULTI PLATFORM

We deem *Telegram* a perfect example of multi platform application:

- 1. A *native* app on all mobile devices
- 2. A *desktop* applications for all platforms (even *FreeBSD*)
- 3. A web interface
- 4. Has a simple and open API to create clients, bots, ...
- 5. Efficiency, security and privacy centered

These traits have continuously inspired the development of *canone3*





Clients – Qt desktop apps

N ×			3	equencer - ι	om:20000 - 1	seneralList		~ ~									
GeneralLi	st Recover SR after BD	~	Start	Stop	Cl	ear		Log	- 1								
Tree Vie	w Node View Recover SR	after BD							÷.								
Device		State	Enable	Block	Executed	Last Exec	Elapsed	Description		_	_	_		a b b		_	_
∽-seq.		FAULT								_	Sec	uencer - 1	om:20000	- GeneralList	.<2>		~ /
>-	eq/check/zerocurrmb_s	OFF		YES	YES	4d 21h 13m	0	check sr current is Zero	Ŧ	Start	Stop	0	lear				Log
	eq/check/rf_s2	OFF	i 🗖 👘	NO	YES	4d 21h 13m	2	Check RF2 OK		Pocovor 6P	ofter PD					soala	
	eg/check/rf_s3	OFF	1	NO	YES	4d 21h 13m	2	Check RF3 OK		necover 3h	alter bD					seq/it	ecover/beamdum
	eg/check/rf s8	OFF		NO	VES	4d 21h 13m	1	Check RES OK			State	Enable	Block	Executed	Last Exec	Elapsed	Description
[]	eq/elicelo/i1_50	011			VEC	44 246 4255	-		np	S	FAULT				unavailable	201	Recover SR af
	ed/cneck/n_s9	OFF		NO	YES	4d 21h 13m		Check RF9 OK	m	nb s	OFF		INO	E YES	4d 21h 16m 4d 21h 16m	2	Check SF CURF
	eq/off/ramp_b	OFF	ļ	NO	YES	4d 21h 13m	3	Switch Booster Injection OFF			OFF	Î.	NO	€ YES	4d 21h 16m	2	Check RF3 OK
	eq/check/mod_p	OFF		NO	YES	4d 21h 13m	2	P active MOD check			OFF	V.	NO	E YES	E 4d 21h 16m	1	Check RF8 OK
>-	eq/check/scw	OFF		NO	YES	-	2	Check SCW ON			OFF	÷	NO	E YES	4d 21h 16m	3	Switch Booste
	eg/check/3hc	OFF	โด	NO	NO	unavailable	0	Check 3HC state			OFF	V	NO	■ YES	⊑4d 21h 16m	2	P active MOD
	ea/check/ns_booster	OFF		NO	VES	4d 21h 13m	8	Check current and state of BOOSTER			OFF	<u> </u>	NO	E YES	E -	2	Check SCW ON
1 C	Leag/check/psc p				1	-41211115111	1	Check catting and state of power su	ste	er	OFF	v	NO	E YES	E 4d 21h 15m	8	Check current
	sed/cneck/psc_p	OFF	{			-		check setting and state of power su	р		OFF	7		~	-	1	Check setting
	>-seq/check/pstcbnch_p	OFF	ł			-	0	Check setting and state of power su	br	ich p	OFF				-	0	Check setting
	>-seq/check/psq_p	OFF	Į			-	0	Check setting and state of power su	ns	p	OFF	┥			-	0	Check setting
	>_seq/check/pslens_p	OFF	Į			-	0	Check setting and state of power sum	p	tb	OFF	Į –			-	<u>ļo</u>	Check setting
	-cog/chock/pch_pth	055	1			1	0	Chack softing and state of power su	<u>p</u>	tb	OFF					0	Check setting
									n M		OFF	┥			-	2	Check setting
efreshed	a subset of 0 sources over 123	7 in 6 millisecond	IS					http 🗸	1 👖		OFF	7			-	0	Check setting
								Iseq/check/	psp o		OFF					0	Check setting
								Iseq/cneck	nsa bi	ts	OFF	$\left\{ \right.$			-	0	Check setting
								seq/check/	psc bt	ts2	OFF	۲.			-	10	Check setting
								seq/check/	psq b	ts2	OFF	7			-	ĺŌ	Check setting
								seq/check/	psb b	ts2	OFF	5			4d 21h 15m	ĵo	Check setting
								seq/off/tunefb)	OFF	\checkmark	NO	■ YES	🛯 4d 7h 35m	14	Switch off Tu
								Log/off/moanfl	<u> </u>				INO	FIVEC	FIAd 7h 25m	11	Switch OEE M
				L				Refreshed a subset of 0	source	s over 1240	in 9 millisecc	nds			(cumbia 👻	Info
	CUM	nnia.	- NT'	rn													

* the same app is *designed* and <u>run</u> transparently regardless the engine in use (native Tango or http/SSE)

- * No engine-specific coding
- * Run with the same command line!
- * A virtual machine run at home can reproduce exactly the control room desktop





Clients – Control Room apps

Approaches to running control room applications remotely

Now

X Control room virtual machines run on the server side

* Although optimized, a full graphical session is streamed, with a strong impact on *bandwidth* and *battery life* of portable devices. If the former today may not be considered a critical limit (however, estimate the load of a server streaming to hundreds of clients), the latter is indeed a precious resource for portable devices and laptops. *The only relevant piece of information is data*.

Then

× Control room virtual machines run on the client

***** The user runs only the apps he needs

* The apps run natively (i.e. the fastest possible on the client device), look exactly as in the control room (they are the same) without lags

× Only data through the network, for native and web apps alike





Clients – Control Room apps

Approaches to running control room applications remotely (II)

Additionally

× Either *virtual* or *physical* machines do not need Tango + dependencies *×* They need only a web browser and Qt for native apps *×* Events only

- * Tango control system potentially *highly relieved* (remember: N clients reading
- *x*, 1 read to the control system)
 - + all benefits discussed in the *design rationale* section





SECTION III

WEB CLIENTS

JavaScript Decentralization





Does ES6 make JavaScript frameworks obsolete?

byTheodoros 'Theo' Karasavvas

https://stackoverflow.blog/2021/11/10/does-es6make-javascript-frameworks-obsolete/





WEB INTERFACES

"Every time JavaScript goes through a major update, we seem to repeat the same cycle. At first, developers are delighted by the new features. They move back to coding directly in JavaScript, and frameworks become less popular. Then, in the relatively long periods between releases, frameworks begin to offer new features and tempt developers back. Repeat."

[React] "building a complex stack of components that you then have to constantly maintain and manage"











"In short, ES6 brought a host of syntactic changes to JavaScript that greatly reduce the need for most frameworks. Coupled with the fact that most of the frameworks we are using at the moment obscure JavaScript code and add an additional dependency, we could see a tangible and permanent reduction in the use of frameworks over the next few years. Or perhaps the cycle will simply repeat, and we'll have only a few years of learning how to write better JavaScript before we retreat into our frameworks again."





2. Decentralization

 We started with a centralized python server from Tango to socket and a PHP server from socket to web (from 2006 to 2017)



- than we moved to a C++ server from Tango or Epics to web with websocket (2017 - 2020)
- NGINX NChan server + SSE (2020 2022)





We changed again Nchan based server to make it more robust, modular and distributed (since 2022)







An other possible configuration could be a web server on any "intelligent" device, for example a websocket server has been implemented on a beaglebone







https://pwm	a.elettra.eu/cordova/	× +											
$\leftarrow \ \rightarrow \ G$	🔒 pwma.elettra.e	eu/cordova/psa_ram	p.php?mod_num=	=2&ramp_name=22.0kV_	7Hz_normal_20	19.ramp						☆	Θ:
🎆 MO	D 2 🛥					22.0k	V_7Hz	z_noi	mal_2	2019.	ramp		A
Open ram	מו:			~	23								
		2010 -			21								
22.0KV_7	Hz_normal	_2019.ramp	Save Sav	e as	20					\checkmark			- 1
t [s]	∆t [s]	V [kV]	$\Delta V [kV]$	$\Delta V/\Delta t$ [V/min]	19				/				- 1
0	0	5	5.00		18	1270 • volta	age ramp:	16.00					
10	10	5.01	0.01	60.0	16		1		-				- 1
490	480	12.0	7.00	875.0	≥ 14		/						
1270	780	16.0	4.00	307.7	Voltag								- 1
2350	1080	19.0	3.00	166.7	12								
3550	1200	21.0	2.00	100.0	10	1							- 1
4750	1200	22.0	1.00	50.0	8								
H					7								- 1
					6 5								
					4		0	1800		700	2600	4500	- 1
					0	90		1000	time [s]	/00	2000	4500	
							-•	- volta	ge ramp				
4												Uishcharte cor	··· · ·

youtube.com/watch?v=9TI2S3mThmQ







youtube.com/watch?v=psj4ZOz7ThA







youtube.com/watch?v=vfe_YrOJvtA







youtube.com/watch?v=I-3L0CaJ5X8





Gun Mod1 Mod2 PS Foc Kly 1 Foc Kly 2 Delta Mod 1 Delta Mod 2

GUN - REAL TIME





youtube.com/watch?v=9TI2S3mThmQ







youtube.com/watch?v=z7FUDB7w2aw





Thank you!







www.elettra.eu