WEB SERVICES AT THE ANF

Antelope Users Group 2012

Reno, NV

OVERVIEW OF SERVICES

- The ANF runs several web sites
- http://anf.ucsd.edu is the primary site for TA and most hosted projects (GLISN, Chile)
- http://eqinfo.ucsd.edu is primarily for the ANZA seismic network, the San Jacinto Fault Zone temporary deployment, and historic projects

RECENT EQS

- Two versions in use
  - Old dbrecenteqs version
  - Interactive version using the Google Maps API

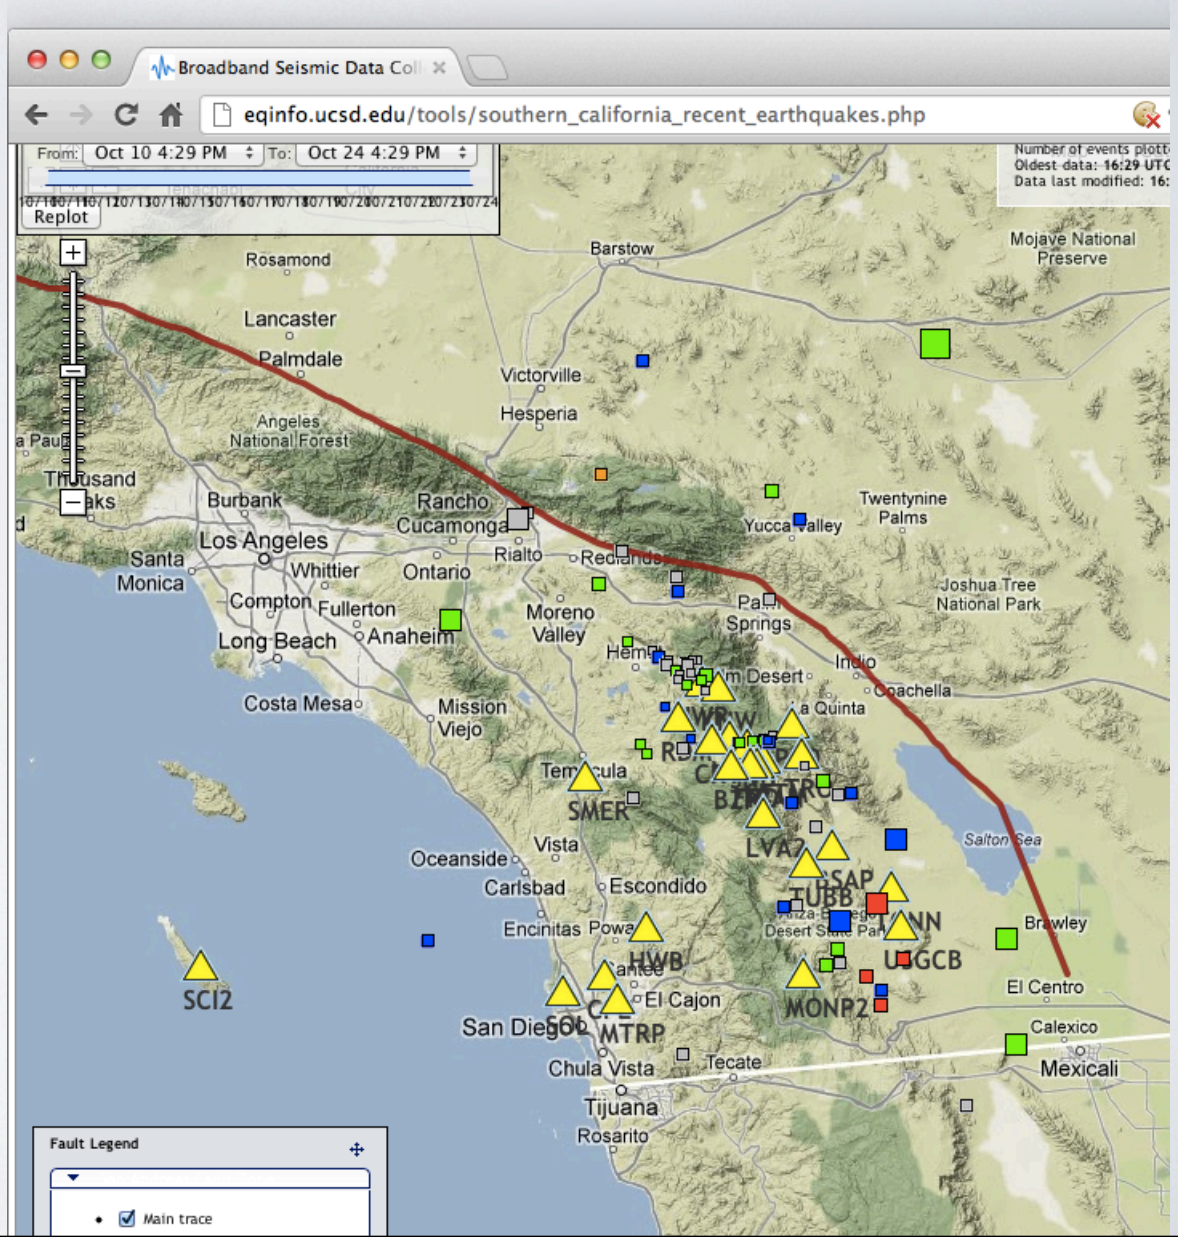

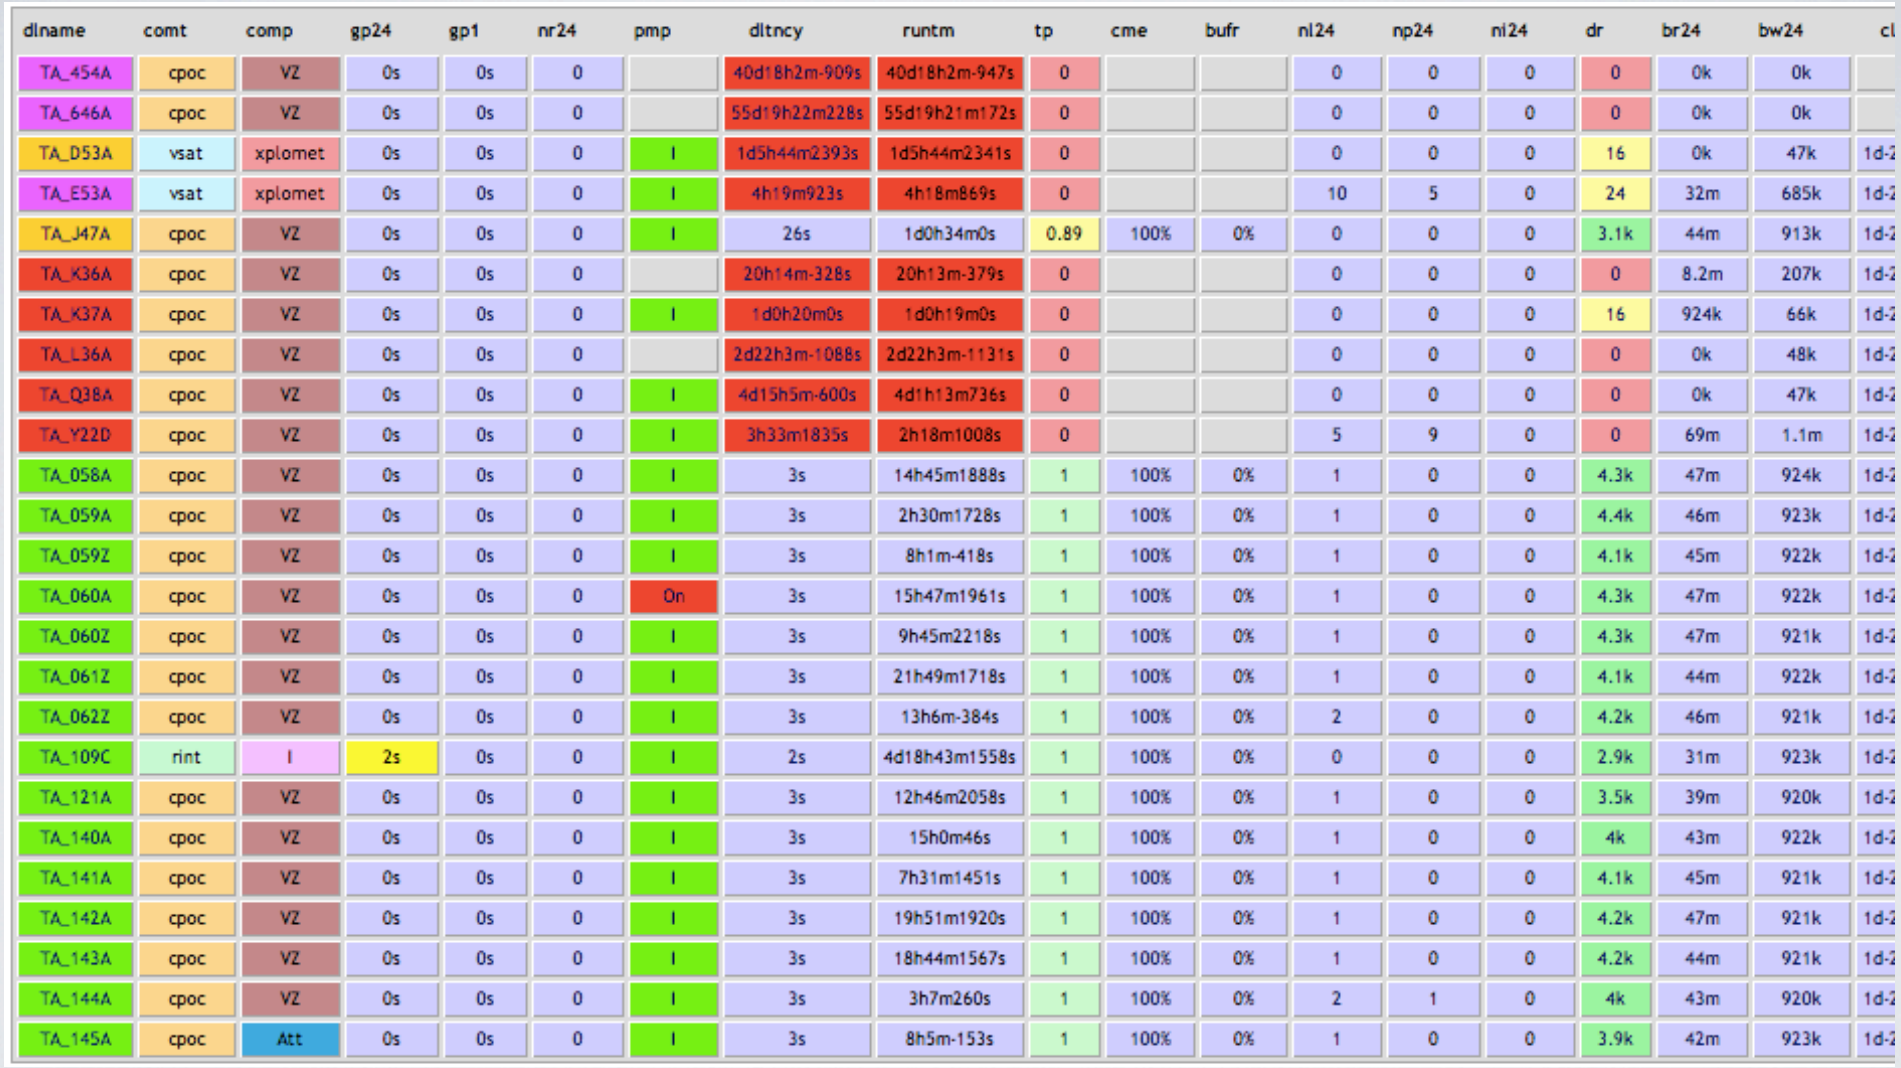

dlname	comt	comp	gp24	gp1	nr24	pmp	dltncy	runtm	tp	cme	bufr	nl24	np24	nl24	dr	br24	bw24	cl
TA_454A	cpoc	VZ	0s	0s	0		40d18h2m-909s	40d18h2m-947s	0			0	0	0	0	0k	0k	
TA_646A	cpoc	VZ	0s	0s	0		55d19h22m228s	55d19h21m172s	0			0	0	0	0	0k	0k	
TA_D53A	vsat	xplomet	0s	0s	0	I	1d5h44m2393s	1d5h44m2341s	0			0	0	0	16	0k	47k	1d-2
TA_E53A	vsat	xplomet	0s	0s	0	I	4h19m923s	4h18m869s	0			10	5	0	24	32m	685k	1d-2
TA_J47A	cpoc	VZ	0s	0s	0	I	26s	1d0h34m0s	0.89	100%	0%	0	0	0	3.1k	44m	913k	1d-2
TA_K36A	cpoc	VZ	0s	0s	0		20h14m-328s	20h13m-379s	0			0	0	0	0	8.2m	207k	1d-2
TA_K37A	cpoc	VZ	0s	0s	0	I	1d0h20m0s	1d0h19m0s	0			0	0	0	16	924k	66k	1d-2
TA_L36A	cpoc	VZ	0s	0s	0		2d22h3m-1088s	2d22h3m-1131s	0			0	0	0	0	0k	48k	1d-2
TA_Q38A	cpoc	VZ	0s	0s	0	I	4d15h5m-600s	4d1h13m736s	0			0	0	0	0	0k	47k	1d-2
TA_Y22D	cpoc	VZ	0s	0s	0	I	3h33m1835s	2h18m1008s	0			5	9	0	0	69m	1.1m	1d-2
TA_058A	cpoc	VZ	0s	0s	0	I	3s	14h45m1888s	1	100%	0%	1	0	0	4.3k	47m	924k	1d-2
TA_059A	cpoc	VZ	0s	0s	0	I	3s	2h30m1728s	1	100%	0%	1	0	0	4.4k	46m	923k	1d-2
TA_059Z	cpoc	VZ	0s	0s	0	I	3s	8h1m-418s	1	100%	0%	1	0	0	4.1k	45m	922k	1d-2
TA_060A	cpoc	VZ	0s	0s	0	On	3s	15h47m1961s	1	100%	0%	1	0	0	4.3k	47m	922k	1d-2
TA_060Z	cpoc	VZ	0s	0s	0	I	3s	9h45m2218s	1	100%	0%	1	0	0	4.3k	47m	921k	1d-2
TA_061Z	cpoc	VZ	0s	0s	0	I	3s	21h49m1718s	1	100%	0%	1	0	0	4.1k	44m	922k	1d-2
TA_062Z	cpoc	VZ	0s	0s	0	I	3s	13h6m-384s	1	100%	0%	2	0	0	4.2k	46m	921k	1d-2
TA_109C	rint	I	2s	0s	0	I	2s	4d18h43m1558s	1	100%	0%	0	0	0	2.9k	31m	923k	1d-2
TA_121A	cpoc	VZ	0s	0s	0	I	3s	12h46m2058s	1	100%	0%	1	0	0	3.5k	39m	920k	1d-2
TA_140A	cpoc	VZ	0s	0s	0	I	3s	15h0m46s	1	100%	0%	1	0	0	4k	43m	922k	1d-2
TA_141A	cpoc	VZ	0s	0s	0	I	3s	7h31m1451s	1	100%	0%	1	0	0	4.1k	45m	921k	1d-2
TA_142A	cpoc	VZ	0s	0s	0	I	3s	19h51m1920s	1	100%	0%	1	0	0	4.2k	47m	921k	1d-2
TA_143A	cpoc	VZ	0s	0s	0	I	3s	18h44m1567s	1	100%	0%	1	0	0	4.2k	44m	921k	1d-2
TA_144A	cpoc	VZ	0s	0s	0	I	3s	3h7m260s	1	100%	0%	2	1	0	4k	43m	920k	1d-2
TA_145A	cpoc	Att	0s	0s	0	I	3s	8h5m-153s	1	100%	0%	1	0	0	3.9k	42m	923k	1d-2

WEBDLMON

SOH PLOTS

- Datalogger state of health metrics
- Communications stats from Cell Modems and VSATS
- Plots last dail, week, month, year, lifetime
- Fed by RRDTool archives

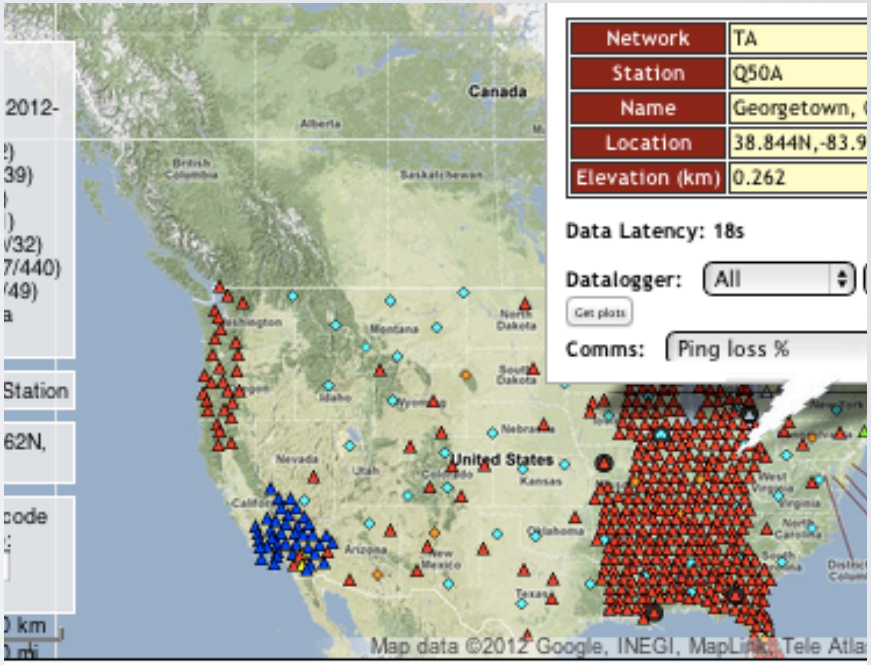

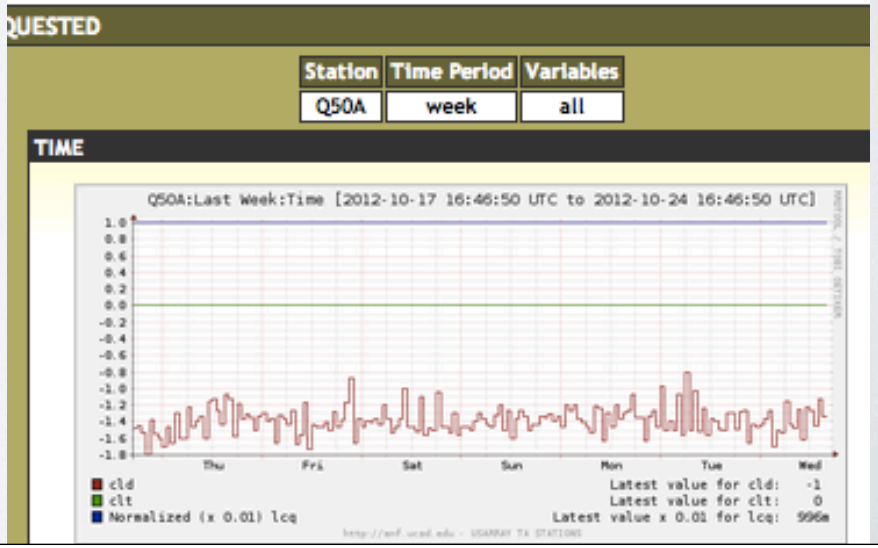

Network	TA - EarthScope Transportable Array Seismic Network				
Station	N47A	Urbana, IN, USA			
Location	40.88 N, -85.69 E				
Elevation	0.25 km				
USArray Ondate	2012-09-04		USArray Offdate	N/A	
Equipment Installation Date	2012-09-04 00:00:00		Equipment Removal Date	N/A	
ANF Certification Date	2012-09-07 20:00:00		ANF Decommission Date	N/A	

Communications	Type	Provider	Power	Duty cycle
	Cellular Modem	Verizon Wireless	N/A	N/A

Instrumentation	Datalogger	Sensor	Infrasound
	Quanterra Q330 [ID tag #: 751]	Guralp CMG-3T [Serial #: T3P05]	MEMS, SETRA, NCPA

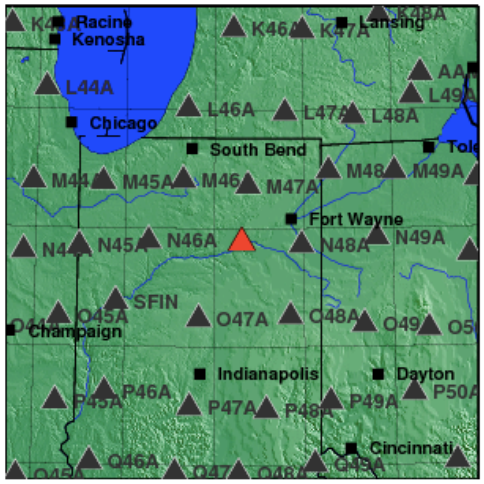

Station history

Comms	Type	Provider	Power	Duty cycle	Ondate	Offdate
	Cellular Modem	Verizon Wireless	N/A	N/A	2012-09-04 00:00:00	N/A

Instrument	Datalogger Type	ID tag	Sensor Type	ID	Ondate	Offdate	Chan.	Hang	Vang
	Quanterra Q330	751	Guralp CMG-3T	T3P05	2012248	N/A	BHE	90	90
	Quanterra Q330	751	Guralp CMG-3T	T3P05	2012248	N/A	BHN	0	90
	Quanterra Q330	751	Guralp CMG-3T	T3P05	2012248	N/A	BHZ	0	0

Infrasound	Sensors Installed	Channel Codes	Ondate	Offdate
	NCPA Infrasound Microphone	BDF_EP	2012248	N/A
	SETRA Absolute Microbarometer	BDO_EP	2012248	N/A
	NCPA Infrasound Microphone	LDF_EP	2012248	N/A
	MEMS Barometric Pressure Gauge	LDM_EP	2012248	N/A
	SETRA Absolute Microbarometer	LDO_EP	2012248	N/A

STATION DETAILS

OTHER IN-HOUSE TOOLS

- Orbmonrtd Image Dumps
- PDF mode graphs (grabbed from IRIS)
- Data return rates
- Interactive waveform explorer (webdlmon)
- Per stations event plots
- Instrument Response plots
- Instrument history charts and plots (useful to see where a datalogger has been deployed before)
- Station Calibration runner

THIRD PARTY WEB TOOLS

- Confluence - Wiki
- JIRA - Issue Tracking and Project Management
- Crowd - Identity Management
- Network Monitoring - Intermapper
- Flikr - online photo archive for station photos
- Github - source code management
- Jenkins - continuous integration

ORIGINAL SITE ARCHITECTURE

- Used the PHP bindings for Antelope written by Kent Lindquist
- All real-time queries to orbs and databases
- Little to no caching of database queries, image generation, etc

CURRENT SITE

- Split web functionality into front-end and back end components
- Lots of back end processing - 3 systems
- Only a single front end, but with some changes we can have multiple front ends

BACK END

- Driven by a bunch of cron jobs and daemons from numerous rtexec instances across three separate systems
- Written in at least 6 programming languages (Perl, PHP, Python, MATLAB, XSLT, Shell)
- A bunch of intermediate products - XML, JSON, Images, Postscript files, etc.

CURRENT FRONT END

- Mixture of PHP (server side) and Javascript (runs on the client in browser)
- Variety of data sources, none directly through Antelope bindings
  - PHP code loads pre-computed XML or JSON
  - Javascript client code loads JSON

GOALS

- Clean up the back end processing
- Reduce the number of languages in use
- Reuse code across networks more effectively

PLAN

- Make web accessible APIs powered by a back end services provider (like Twisted)
- Expand upon existing JSON feeds
- Convert static JSON files to dynamic feeds
- Use APIs as building blocks for pages
- Use a web framework (MVC) for the front-end display

BACK-END CONSIDERATIONS

- Some things should be "cron jobs"
- Data changes very infrequently
  - Database queries of station metadata
  - Other db queries - things driven by the dbmaster and dbops

BACK-END CONSIDERATIONS

- Other things should be quasi real-time
  - Orb queries
  - Datalogger status packets - current value and near term historic graphs
  - Real-time waveform streaming

API EXAMPLES

- Preliminary - don't depend on these being in our final version
- GET /nets/TA/stations?status=active
- GET /dlstatus/TA/stations/TA_109C
- GET /dlstatus/TAprelim/stations
- GET /nets/TA/dataloggers/109C

FRONT END STRATEGY

- Initially migrate existing pages from static text files to API-provided data
- Migrate to a web framework. Use templating where possible.
- Re-usable modules for displaying station details, real-time data logger monitoring, SOH plots

WHERE WE ARE NOW

- Working on code for a real-time feed of datalogger status -- replacement for orbdlstat2xml
- Written using Python and the Twisted framework
  - Queries multiple orbs and can consolidate them into the same feed or multiple feeds
  - Orb queries are asynchronous - web queries are not affected while new data is loaded
- Rewriting webdlmon as a consumer for data provided by this daemon

FUTURE SERVICE IDEAS

- orbwf2decimatedJSON
  - Could power orbmonrtd for a web browser
  - reads waveform data from the orb
  - pre-decimates the data for serving up to web browser clients

FUTURE SERVICE IDEAS

- Rewrite special events pages to be database driven.
- Currently a static mix of data and formatting code
  - Breaks when we change the look and feel of the site

CURRENT STUMBLING BLOCKS

- Twisted works best with a deferred thread model, but Antelope bindings do not release the global interpreter lock
  - Worked around with Python Ctypes to call C libraries directly
- Using post-release 5.2-64p since we have problems with the current python bindings

CURRENT STUMBLING BLOCKS

- Database reads of waveforms from miniseed.
- Continuous updates to blockettes cause problems with trexcerpt since blockettes change during trexcerpt call.
  - Goes back to fundamental limitations of Datascope IPC - it's confined to a single server.

CURRENT STUMBLING BLOCKS

- Python libraries are lacking exceptions. Instead you are using C style return value checking.
  - Elog exceptions in prior versions weren't "Pythonic" - they were too general, and didn't reflect the actual type of error that occurred.
  - Would like a python-specific set of exceptions