



THE ADVENTURES OF AARNET ACROSS THE EOS DIMENSION

David Jericho, Systems Architect

5th February 2018

DAVID JERICHO

Systems Architect

20 years experience

13 years at AARNet

...and still not bored

Motorcycle crazy

Ice hockey player

Wanted to be a pilot

Started coding at 5

Parent of 3 children

Likes his dark office

Forgets he can code

Still owns SGI kit



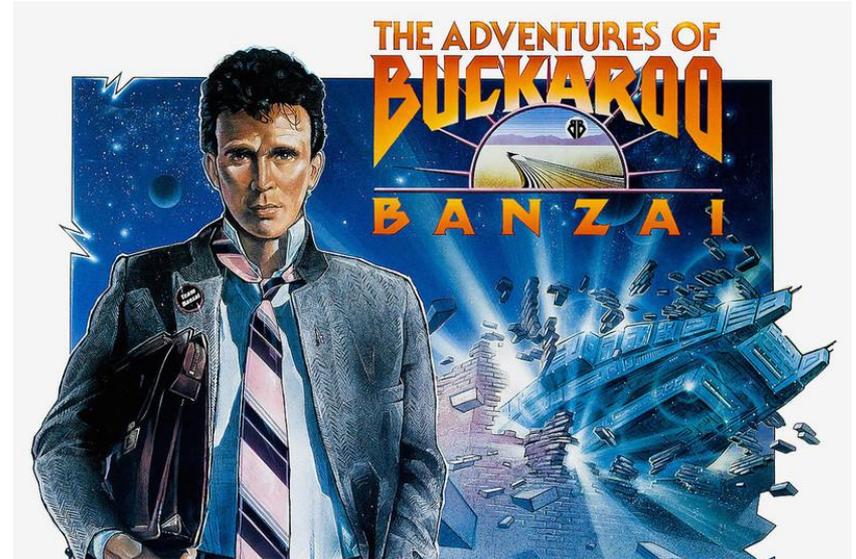
This topic reminded me of the 1984 movie *“The Adventures of Buckaroo Banzai Across the 8th Dimension!”*

WHY THE MOVIE REFERENCE?

“The premise centers upon the efforts of the polymath Dr. Buckaroo Banzai, a physicist, neurosurgeon, test pilot, and rock musician, to save the world...”

Many very talented people producing work that benefits others, regardless of people understanding what that work is.

It wouldn't surprise me if someone here was in a rock band.



https://en.wikipedia.org/wiki/The_Adventures_of_Buckaroo_Banzai_Across_the_8th_Dimension

SO, AARNET NEEDED SOME STORAGE...



We tried

- GlusterFS
- XtremFS
- Lustre
- Lizard/MooseFS
- ...many others...
- And, Ceph

WHY CEPH WASN'T THE ANSWER

- It doesn't handle network latency
- The filesystem paradigm still isn't ready
- Isn't yet truly scaling to the tens of PB
 - Sage Weil (Barcelona; 2016) points out OSD count has only been tested to 12,000
 - No massive production clusters people are sharing information on
- At the risk of creating a false cause fallacy:
 - If it did work at even a fraction of exa-scale, the big sites wouldn't be busy writing their own filesystems



EOS IS THE ANSWER FOR AARNET



- It works at continental scale latency
- It provides a familiar paradigm for access
- It's designed for strong integrity
- We talk to the primary authors and users daily
 - This has matured greatly over time
- It scales very well to match our lofty goals

ISSUES ALONG THE WAY

- Bugs in code
- Missing features that were in documentation
- No documentation for features
- Awareness of the whole CI/CD/source chain
- Advance knowledge of roadmap
- Communication with the developers



MANY OF THESE HAVE BEEN IMPROVED



- Bugs in code
 - CERN is always responsive
- Missing features that were in documentation
 - These are fixed once pointed out
- No documentation for features
 - Once again, fixed when raised
- Awareness of the whole CI/CD/source chain
 - CERN has made much of this available
- Advance knowledge of roadmap
 - CERN always shares information when we visit
- Communication with the developers
 - It's fun chatting with many of them on Slack

NOTICE A PATTERN?

- In every case, we depended upon CERN
- Don't misunderstand, CERN have been excellent
- And we think you guys are rock stars
- However only 37 of the 80 attendees on the participant list report themselves as being from CERN itself
- The community is growing and spreading



WALKING WITH MASSIMO



- On many of the opportunities for one-on-one conversations with Massimo, he has often said two things:
 - There is no need to reinvent the wheel
 - Groups like CERN reached this scale before others, so use their knowledge to solve issues
- The corollary here is:
 - As the EOS community has grown, it's repeating patterns of open source projects, and we should adapt and learn from successful efforts



An idea to consider

HALLMARKS OF SUCCESSFUL OPEN SOURCE PROJECTS

Talking about successful open source projects, Linux Foundation Executive Director Jim Zemlin mentions important “janitor” functions, such as:

- triaging bugs
- answering questions from new users and new developers
- dealing with trademark and license issues
- and generally being available to grease the inevitable frictions that occur in any kind of large collaboration.



<https://www.linuxfoundation.org/blog/successful-open-source-projects-common/>

EXPANDING ON JANITOR FUNCTIONS



<https://dzone.com/articles/running-a-successful-open-source-project-1>

According to Wayne Beaton (Eclipse foundation):

- Clear definitions of success
- Documentation must be clear and up to date
- Roadmaps must be public and findable
- Consider a hosting foundation
- Be agnostic about the source of contribution
- Strong brand and IP management
- Communication must be participated in

- Have a plan!

CEDING CONTROL?

No!

Merit continues to justify participation, and..

The same set of rules apply to everyone...

Not to mention the existing licenses already clearly state what and where the work came from...

As is often the case, the project originator continues being the project lead.



CLOSING QUESTIONS



Is it time to consider the future with a foundation or focused lead agent behind the open source effort?

How, and what, or even if?



THANK YOU
