

Sysadministrivia

Linux, Lagers, and Late Nights

S4E10: "Unavailable"

Posted 2019-07-08 03:59

Modified 2019-07-07 06:06

Comments 2

Navigation

Previous Episode	Next Episode
S4E9: "Great Walls of Fire"	S4E11: "SCADA isn't an STI"

Log

Recorded (UTC)	Aired (UTC)	Editor
2019-06-27 02:36:29	2019-07-07 06:02:38	"Edita"

Verification

Format	SHA256	GPG	Audio File
MP3	796b0988c7e3e342b7199b5033c3a625559bbb6b21ebcbf64a5e8c0f499164db	click	click
OGG	098d6e52c7636fe6c27d7a063dfc25a50f5a01e3910d2fb4f1e88894b8e52bfc	click	click

Quicklisten:

We talk about load balancing, and just how impossible it is to do right.

- Just the Tip
- Notes
- 15 Clams
- Errata
- Music

Just the Tip

- Paden gives an introduction to RAID and a brief rundown of its various applications.
 - The Wikipedia article for RAID levels
 - mdadm is extremely useful.
 - LVM is not real RAID. It's not even software RAID. I call it "virtual RAID" because to the OS it can behave **like** RAID but is not in itself any sort of actual RAID implementation.
 - Same with ZFS "RAID" (pools) – not real RAID.

Notes

Starts at **21m45s**.

I was drinking water. Paden was drinking Absolut vodka and water. Jthan was drinking Busch.

- HA (high-availability)
 - If you need something **absolutely, positively** available, in the case of losing millions of potential dollars or more for each second it's unavailable.
 - HA is **EXCEEDINGLY** expensive. The cost increases exponentially the more you want it to be available, and you very, very quickly reach the point of diminishing returns.
 - 100% availability is literally impossible, because physics and probability.
 - The absolute golden rule when it comes to HA: "If you need high-availability, you can afford it. If you can't afford it, you don't need it."
 - This works because of scaling – bigger companies deal with bigger amounts of money, and therefore have a greater sensitivity to the need for HA.
- What **isn't** HA:
 - Round-robin DNS (it's absolutely pointless)
 - Among the countless negative things to it, the primary point is it's entirely unpredictable (namely due to how DNS resolution occurs).
 - Some resolvers don't even randomly pick a record for round-robin names; they just pick the first one specified, etc.
 - Load balancing (it's a *component* of HA but not the entirety)
- What **is** HA:
 - True high-availability **requires**:
 - Fault-tolerancy at the network, hardware, physical, and political level
 - Load balancing (in whichever methodology makes sense for the application) as a component for the HA
 - **Immediate** failover and cold/hot swapoff

- Anycast DNS *and* geocast/anycast IP addressing
- Hardware HA, unless you manually switch out nodes for maintenance, need to have hot-swappable hardware (disks, RAM, CPUs, NICs, etc.)
- Why is complete availability impossible?
 - Every single point of balance has at least one point of failure.
- Scope
 - **“Realistically**, for my business, what is an attainable level of availability?”
- Fake it ‘til you make it:
 - HAproxy and/or Nginx (limited support in the open-source/“community” version and it requires a fairly recent version) can let you get **close** - around 3-4x “nines”.
 - There are also some commercial products, like F5 Networks (whom recently bought Nginx).
 - And for the network stack itself, IPVS/the LVS project goes a LONG way in making networks more hardy but it requires you to have full control over your network.

15 Clams

In this segment, Jthan shares with you a little slice of life. The title is a reference to this video. (2m16s in)

Starts at **56m03s**.

Jthan didn’t know that you can’t use empty list as default parameter values in functions in Python.

Errata

- We all described RAID levels totally wrong (e.g. RAID5 is distributed parity). Just read the wikipedia article. The performance hits we mentioned should be still accurate, though.
- I, for SOME incomprehensible reason (probably overexcitement), kept saying “H-I” instead of “H-A”.
- Yes, Jthan, Cloudflare DOES have a CDN offering!
- Platter disks are extremely sensitive to noise (careful of the noise!).
- I TOTALLY MISSED A “BLUE OYSTER CULT” JOKE/REFERENCE WHEN JTHAN WAS INTRODUCING THE 15 CLAMS! UGH.
- The “hidden track” at the end is a recording I grabbed from pre-show when Jthan was running to the kitchen. We heard him clanging around, but we caught him getting mad at his cat (Julius), which was icing on the cake. He said the cat wouldn’t leave the room and would scratch at the door to be let out while recording. Press F for Julius.

Music

Music Credits

Track	Title	Artist	Link	Copyright/License
Intro	Shaanti	Rolemusic	click	CC-BY 4.0
Outro	Daybreak	Chad Crouch	click	CC-BY-NC 3.0

(All music is royalty-free, properly licensed for use, used under fair use, or public domain.)

Author r00t^2

Categories Season Four

Comments

1.

2019-07-10 09:32 (744 days ago)

OK, where’s the picture? ;)

2. 2019-07-10 19:58 (743 days ago)

@Jon-

We’re still waiting on amayer’s OK to post it! I asked Jthan to get disclosure but he hasn’t yet.

Also, saw your message in the channel this morning! We don’t actually have +m on the channel, it’s +nrt (https://freenode.net/kb/answer/channelmodes). If anything, I’d guess the +r is causing issues for you but it’s the only thing preventing spam. Have you ensured that you auth’d with Nickserv first when connecting via Matrix?