



# What's new in knet this year? (2020)

Or “You’ve been tinkering with this for ages. It must work by now?”

---

Christine Caulfield

Principal Software Engineer, Red Hat

# What is knet?

A quick overview for those that are new, or simply forgetful

Knet is the network transport for corosync 3

- Multi-home, Multi-transport, Dynamic MTU
- Blah blah blah
- You get the idea
- Provides all the dirty comms stuff
- Encryption, compression, link events, stats
- STATS!



# What is knet? (cont)

A quick overview for those that are new, or missed the last slide

**Knet is the network transport for corosync 3**

- We ripped out LOADS of the corosync code
- Including RRP
- Cos knet does all of that
- We also get dynamic link configuration
  - Which is nice
- V1.0 released Jan 2018
- Current version 1.14



# Any questions so far?

This isn't going to get any easier



# New stuff!

## Features & bug fixes (overview)

### Features first, then you can have a snooze

- ACLs
- libnozzle
- Compression types
- API man pages
- Better thread handling



# New stuff(1)

## ACLs


- Prevents unauthorized access to the cluster
- By default nodes not known to k8s will be ignored
- Yes, you can do this with crypto but this is neater
- And quicker
- We were going to leave this till v2 but Honza pre-empted us
- Corosync only supports on/off
- Using the library, lots of options possible
- A bit like iptables but we won't change it every year
- You hope



# New stuff(2)

## libnozzle

- More useful than it sounds
- Creates an ethernet device that tunnels through knet
- Easier than all that CPG stuff
- No extra coding required to use it, just use IP address
- You get all the transport options & failover for free
- In corosync 3.0.2
- Used to be called libtap
  - But that was taken
  - And nozzle sounds weirder







# New stuff(3)

## More compression types

- Actually we only added libzstd
- But it shows it's nice and easy
- And I wanted to use this photo
  - it was taken from a train ...



# New stuff(4)

## man pages

- A small thing, I admit
- But it shows we care
- All man pages show the API calls and structure definitions used
- Generated from the header file
  - Doxygen
  - + some terrible thing I knocked up one afternoon
  - Seriously who wants to be typing `.TP` for hours on end?



# New stuff(5)

## Better thread handling

- Like you care
- But we do, and every little helps
  
- Improved regression tests
  - To avoid more ...



# Bugfixes

## Yes, there were bugs

- Varying levels of bugs fixed
- From ‘how did this ever work?’
- To typos in comments
  
- A lot of fixes in the config system
- Preventing bad configurations
- Flush on close – helps with corosync LEAVE messages
- Crypto reconfig (not corosync, and not really what you think it is)



# Bug fixes 2

- MTU fixes – lots of these. I told you last time it was a lot of work
- SCTP was a bit of a ... train crash
- All the little stuff – covscans, flying pointers etc etc
  
- Thanks to all those who reported bugs
  - And helped with reproducers
- Seriously, it's great working with you



# CI (I was told to talk about this) - on Monday

- Covers all of the clusterlabs projects
- libqb,knet,corosync,pacemaker,sbd ... all that stuff
- jenkins
- Multiple architectures & versions
- Coverity scans
- [ci.kronosnet.org/view/<project>](http://ci.kronosnet.org/view/<project>)



[Back to Dashboard](#)[Status](#)[Changes](#)[Workspace](#)[Build with Parameters](#)[Delete Multi-configuration project](#)[Configure](#)[GitHub](#)[Rename](#)

## Project corosync-manual-build-all

[add description](#)[Disable Project](#)

### Configurations

[centos-7-x86-64](#) [debian-experimental-x86-64](#) [debian-testing-x86-64](#) [debian-unstable-sparc64](#) [debian-unstable-x86-64](#) [debian10-x86-64](#) [debian9-x86-64](#) [fedora-rawhide-x86-64](#) [fedora30-x86-64](#) [fedora31-x86-64](#) [freebsd-12-i386](#) [freebsd-12-x86-64](#) [freebsd-devel-x86-64](#) [opensuse-15-x86-64](#) [opensuse-tumbleweed-x86-64](#) [rhel75z-power8-ppc64le](#) [rhel75z-s390x](#) [rhel75z-x86-64](#) [rhel76z-power8-ppc64le](#) [rhel76z-s390x](#) [rhel76z-x86-64](#) [rhel77z-power8-ppc64le](#) [rhel77z-s390x](#) [rhel77z-x86-64](#) [rhel81z-power8-ppc64le](#) [rhel81z-s390x](#) [rhel81z-x86-64](#) [ubuntu-18-04-lts-i386](#) [ubuntu-18-04-lts-x86-64](#) [ubuntu-19-04-i386](#) [ubuntu-19-04-x86-64](#) [ubuntu-19-10-x86-64](#) [ubuntu-devel-i386](#) [ubuntu-devel-x86-64](#)

### Permalinks

- [Last build \(#50\), 1 mo 27 days ago](#)
- [Last stable build \(#50\), 1 mo 27 days ago](#)
- [Last successful build \(#50\), 1 mo 27 days ago](#)
- [Last failed build \(#45\), 2 mo 14 days ago](#)
- [Last unsuccessful build \(#48\), 1 mo 27 days ago](#)
- [Last completed build \(#50\), 1 mo 27 days ago](#)

### Build History

[trend](#)


find

<a href="#">#50</a>	8 Dec 2019, 03:04
<a href="#">#49</a>	8 Dec 2019, 00:41
<a href="#">#48</a>	8 Dec 2019, 00:36
<a href="#">#47</a>	21 Nov 2019, 04:39
<a href="#">#46</a>	21 Nov 2019, 04:34
<a href="#">#45</a>	21 Nov 2019, 02:14
<a href="#">#44</a>	21 Nov 2019, 00:59
<a href="#">#43</a>	21 Nov 2019, 00:56
<a href="#">#42</a>	21 Nov 2019, 00:55
<a href="#">#41</a>	21 Nov 2019, 00:52

# Future

## Stuff for next time

- Proper crypto config changes while running
  - Making good progress on this
- Re-keying
- **More stats?**
  - Yes really
- We'll think of other stuff



Train photo (c) railcam.uk



# RealttimeTrains.co.uk

Running on corosync

RTI NEWS SEARCH CONTACT SHOP Quick Search GO

## 0627 MANCHESTER PICCADILLY TO SCARBOROUGH

DEPARTING ON 3RD FEBRUARY 2020

OPERATED BY  
**TRANSPENNINE EXPRESS**

- First & Standard class seating
- Reservations available

Location	PI	Planned		Realtime		
		Arr	Dep	Arr	Dep	Dly
MAN Manchester Piccadilly	1		0627		<b>0627</b>	
SYB Stalybridge	4	0641	0642	<b>0641</b>	<b>0643</b>	+1
HUD Huddersfield	8	0659	0700	<b>0700</b>	<b>0702</b>	+2
LDS Leeds	15	0722	0727	<b>0721</b>	<b>0727</b>	
GRF Garforth	1	0735	0736	<b>0735</b>	<b>0737</b>	+1
YRK York	5	0758	0803	<b>0755</b>	<b>0803</b>	
MLT Malton		0827	0828	<b>0827</b>	<b>0828</b>	
SEM Seamer		0844	0845	<b>0845</b>	<b>0846</b>	+1
SCA Scarborough	3	0852		<b>0851</b>		-1

Go to detailed version

# Links 1

Website <https://kronosnet.org/>

- Tarballs <https://www.kronosnet.org/releases/>
- Source code <https://github.com/kronosnet/kronosnet/>
- 

• TODO list <https://trello.com/kronosnet>

• Google docs <https://goo.gl/9ZvkLS>

• IRC: **#kronosnet** on Freenode

•

Mailing lists

• <https://lists.kronosnet.org/mailman/listinfo/users>

• <https://lists.kronosnet.org/mailman/listinfo/devel>

• <https://lists.kronosnet.org/mailman/listinfo/commits>

•



Photo *duesentrieb* on flickr (CC BY-NC-ND 2.0)

# Any Questions?

