

MySQL Ecosystem on ARM

By Krunal Bauskar
Driving #mysqlonarm initiative



about me

- Working in MySQL space for more than decade now.
- In past worked
 - @ Percona as PXC product lead
 - @ Oracle/MySQL as InnoDB Developer
 - @ Yahoo! Labs as Big-Data Research Engineer
 - @ Startup (now part of Teradata) as MySQL Engineer (world fastest storage engine)
 - ... more to add
- Currently working
 - @ Huawei - Open Source DB group driving [#mysqlonarm](#) initiative
 - driving all variants of mysql: mysql, mariadb, percona and ecosystem tools.

agenda

- growing arm ecosystem
- mysql/mariadb/percona on arm
- why run dbs on arm?
- state of ecosystem
- moving forward

agenda

- growing arm ecosystem
- mysql/mariadb/percona on arm
- why run dbs on arm?
- state of ecosystem
- moving forward

growing arm ecosystem

The ARM logo consists of the lowercase letters "arm" in white, centered within a solid blue square.

arm

- cell phones and network equipment

growing arm ecosystem

The ARM logo consists of the lowercase letters "arm" in white, centered within a solid blue rectangular box.

arm

- cell phones and network equipment
- housing appliance & IoT

growing arm ecosystem

The ARM logo consists of the lowercase letters "arm" in white, centered within a solid blue square.

arm

- cell phones and network equipment
- housing appliance & IoT
- automobiles/space/defense

growing arm ecosystem



arm

- cell phones and network equipment
- housing appliance & IoT
- automobiles/space/defense
- high-performance computing

NEW!

growing arm ecosystem



what is fueling hpc growth on arm ?

- lower ownership cost thereby improved cost/performance (go-green initiative 🌱)
- increasing ecosystem
- growing user-base/developer-base
- majority of the data-generating device using arm
- availability of arm instances through cloud
 - huawei-cloud - kunpeng 920
 - aws - graviton 2
 - oracle cloud/packet.com - ampere altra

growing arm ecosystem



- majority of os providers now provide arm port
- applications are being ported (big-data, databases, application server, load-balancer)
- arm cloud instances are estimated to provider 30% saving (compared to x86 instance of same class).
- by 2030* arm instances would be most booted instances for running hpc software

growing arm ecosystem



ARM is there for quite sometime why it is taking that long to port things to ARM?

- Weak Memory Model
- Low Level Code (Timer/SpinLoop)
- More Core/NUMA compatible
- Use of ACLE
- NEON parallelization
- Cache Line-Size/Branching/Pipeline Differences

agenda

- growing arm ecosystem
- mysql/mariadb/percona on arm
- why run dbs on arm?
- state of ecosystem
- moving forward

mysql on arm



- started releasing packages for arm starting 8.x
- no major/dedicated optimization yet (till 8.0.21) except bug fixes.
- active community patches (30+).
- some good traction expected in 8.0.22+ (few already committed. more work-in-progress).
- other ecosystem tools on arm
 - mysql-router (load-balancer)
 - mysql-shell (smart-shell)
 - mysql-connector (connector)
 - group-replication (inherent to server)

mysql on arm



- major contributions around
 - crc32/crc32c (hardware optimization)
 - memory barrier (arm weak memory model)
 - numa-aware constructs (counter/connection)
 - 64-bits optimizations (8 bytes copy, etc..)
 - cache-line optimizations
 - switch to/optimal use of atomics (c+11)
 - spin-loop optimization

percona on arm



- no official packages yet but compiles and run successfully on arm (except tokudb and myrocks)
- fully upstream (mysql) compatible so has all the needed fixes/optimizations (for arm) from upstream.
- enjoy good community support with multiple users/developer tried running percona-server on varied arm architecture.

mariadb on arm



- official packages available on arm (centos, rhel and ubuntu). more distro in progress.
- first to support and optimize on arm.
- actively accept arm patches (including correctness and performance)
- evaluation of feature involves considering performance on arm too.
- lacks ecosystem support (connector/maxscale) for arm.

mariadb on arm



- optimizations (mostly folded).
 - numa scalability
 - spin-loop optimization
 - memory barrier
 - crc32/crc32c
 - port to use atomic (all over)
 - lse/branching/cache line optimization
 - timer-counter optimization

agenda

- growing arm ecosystem
- mysql/mariadb/percona on arm
- why run dbs on arm?
- state of ecosystem
- moving forward

cpm



- arm usp is lower cost and more computing resources
- needed a model to help exploit this fact and also compare it with the current baseline numbers to see if real cost saving is being achieved.
- we developed a cost-performance-model (cpm) to exploit this fact by keeping cost constant thereby allowing more arm resources in turn more tps/usd.

<https://mysqlonarm.github.io/CPM/>

cpm



	Huawei Cloud	AWS	packet.com
ARM	2745 USD (yearly) 16 vCPU/64GB (2.6 Ghz Kunpeng 920 based instance)	3384 USD (yearly) 16 vCPU/64GB (2.5 Ghz Graviton 2 based instance)	8760* USD (yearly) 32 vCPU/128GB (3.3 Ghz Ampere eMAG with storage)
x86	4687 USD (yearly) 16 vCPU/64GB (3.0/3.4 Ghz Intel Xeon 6266C Cascade Lake based instance)	4239 USD (yearly) 16 vCPU/64GB (2.5/3.1 Ghz Intel Xeon 8175M Skylake)	17520 USD (yearly) 32 vCPU/256GB (2.5 Ghz AMD EPYC 7502P with storage pricing)
Saving	41%	20%	~ 50% (with storage)

<https://www.huaweicloud.com/intl/en-us/pricing/index.html?tab=detail#/ecs>

<https://aws.amazon.com/ec2/pricing/on-demand/>

<https://www.packet.com/cloud/servers/>

cpm



Can we leverage the said cost saving and still get improved or on-par throughput?

Challenges with existing model:

- CPU frequency differences. (2.5, 2.6, 3.1, Turbo mode).
- CPU generation differences (supporting advanced instruction set).
- x86 physical cores -vs- HT cores (vCPU as per cloud terminology)
- NUMA arrangements.
- Memory differences.
- Scheduler (especially for more-numa) issues. (improved kernel support).

cpm



Keeping cost constant let's analyze if we can get more throughput (there by more tps/per USD) from ARM based instances

given x usd user can get

- m arm resources
- n x86 resources (where $m > n$)

can these extra arm resources helps improve performance?

Given the on-par cost of both the resources can we exploit the ARM variant to its fullest benefit and get better performance when compared to x86?

cpm



- on huawei-cloud for 3500 USD (yearly)
 - x86: 12 vCPU/48GB - 3432 (3.0 Ghz)
 - arm: 24 vCPU/48GB - 3578 (2.6 Ghz)

almost double the compute power (36 -vs- 62.4)
difference increases with bare-metal (2.5-5x)



- benchmarking

- Server Configuration:

- 100 tables * 3 millions (69 GB)
- 80 GB (CPU Bound) /35 GB (IO Bound)
- 20 GB of redo-log
- MySQL-8.0.21 (latest GA release)
- Percona Server 8.0.21 (release branch)
- MariaDB Server 10.5.6 (latest GA release)

- Test-Scenarios:

- sysbench: point-select, read-only, read-write, update-index, update-non-index

- Machine Configuration (bare-metal)

- **x86_64:** Intel(R) Xeon(R) Gold 6151 CPU @ 3.00GHz (*HT enabled*) [*28 ht-cores: 22 server + 6 client*], 192GB mem
- **ARM:** Kunpeng 920 (2.6 Ghz) [64 cores: 56 server + 8 client], 192GB mem






- Storage

- 1.6TB NVMe SSD (random read/write 180K/70K)

mysql on arm



cpu-bound
uniform
8.0.21






https://github.com/mysqlonarm/benchmark-suites/blob/master/sysbench/conf/mysql.cnf/100tx3m_8x_cpbound.cnf


mysql on arm



cpu-bound
uniform
8.0.21








ARM continue to score with higher scalability in all cases.



mysql on arm



io-bound
uniform
8.0.21







https://github.com/mysqlonarm/benchmark-suites/blob/master/sysbench/conf/mysql.cnf/100tx3m_8x_iobound.cnf


mysql on arm



io-bound
uniform
8.0.21








ARM scores even for IO bound workload that involves heavy flushing.



mysql on arm



cpu-bound
zipfian
8.0.21







https://github.com/mysqlonarm/benchmark-suites/blob/master/sysbench/conf/mysql.cnf/100tx3m_8x_cpbound.cnf


mysql on arm



cpu-bound
zipfian
8.0.21








ARM scores even for higher contention workload.



mysql on arm



io-bound
zipfian
8.0.21







https://github.com/mysqlonarm/benchmark-suites/blob/master/sysbench/conf/mysql.cnf/100tx3m_8x_iobound.cnf


mysql on arm



io-bound
zipfian
8.0.21



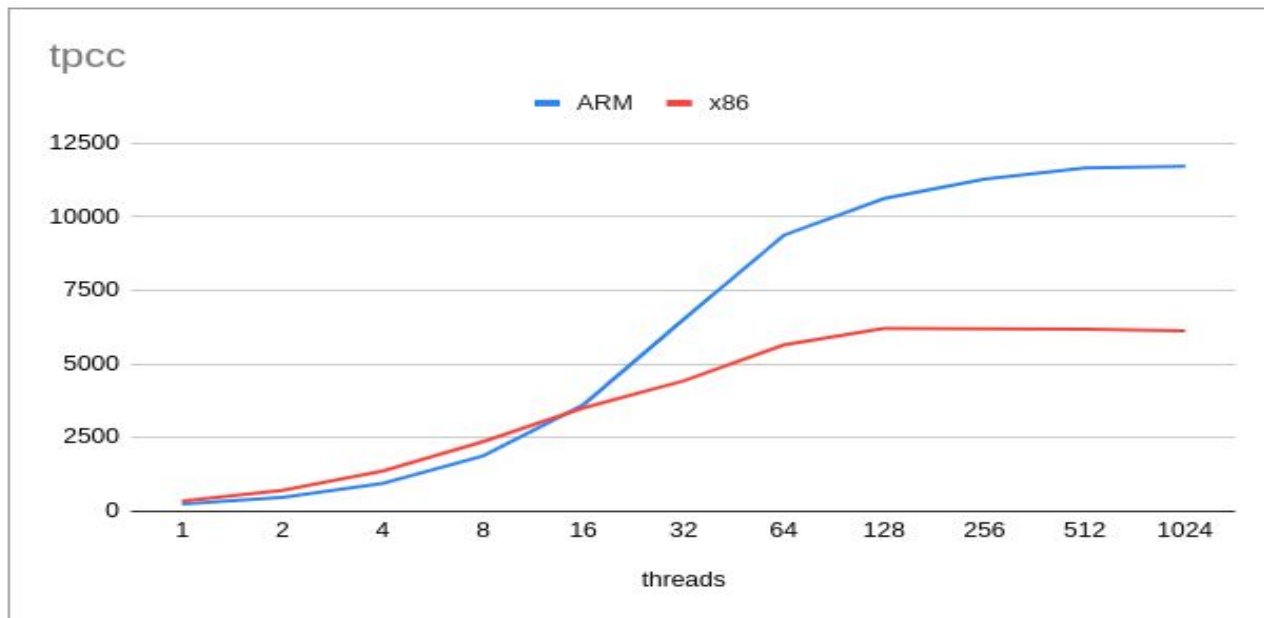
“IO bound higher contention”
...ARM still leads.



mysql on arm



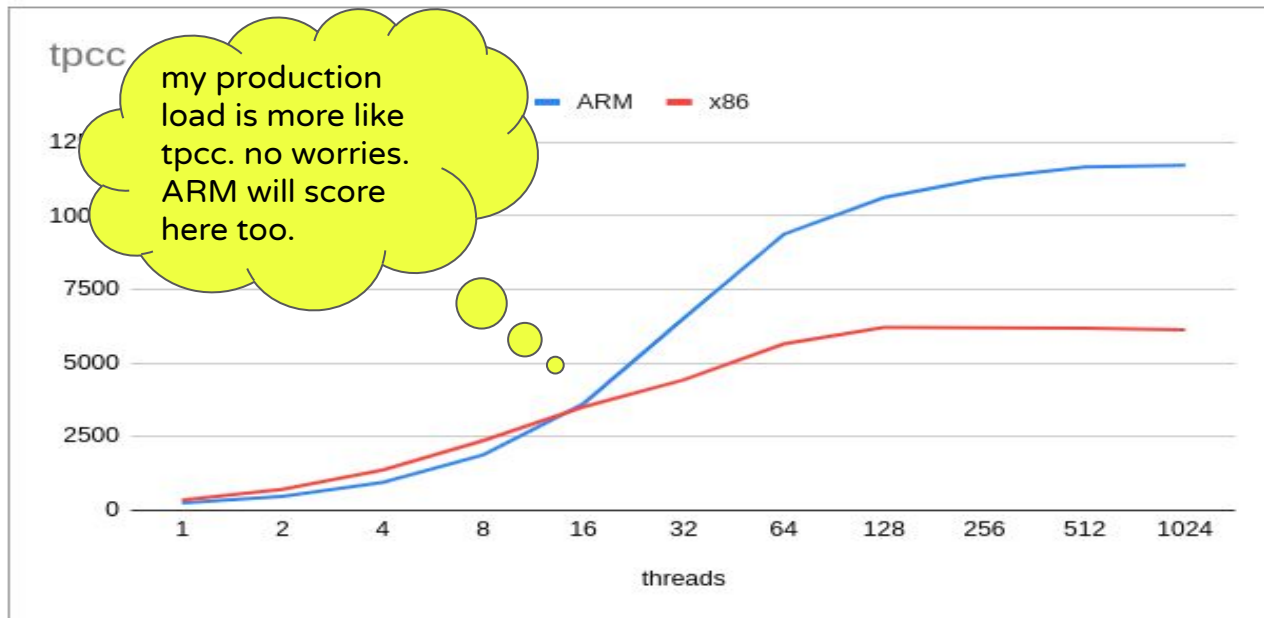
tpcc
cpu-bound
8.0.21




mysql on arm



tpcc
cpu-bound
8.0.21



mysql on arm








percona on arm



PERCONA
SERVER

cpu-bound
uniform
8.0.21








https://github.com/mysqlonarm/benchmark-suites/blob/master/sysbench/conf/percona.cnf/100tx3m_8x_cpubound.cnf

percona on arm



PERCONA
SERVER

io-bound
uniform
8.0.21








percona on arm



PERCONA
SERVER

cpu-bound
zipfian
8.0.21








percona on arm



PERCONA
SERVER

io-bound
zipfian
8.0.21



percona on arm



PERCONA
SERVER

io-bound
zipfian
8.0.21






I can run faster
with ARM



mariadb on arm








cpu-bound
uniform
10.5.6



mariadb on arm



cpu-bound
zipfian
10.5.6




mariadb on arm



MariaDB

cpu-bound
zipfian
10.5.6



why run db on arm?

The logo for CPM (Cloud Performance Monitor) features the letters 'CPM' in a bold, gold-colored, sans-serif font. The letters are set against a dark, textured rectangular background that resembles a carbon fiber or similar material.

what does increase performance means?

why run db on arm?

The logo for CPM (Cost Per Million) is displayed in a dark, textured rectangular box. The letters 'CPM' are rendered in a bold, gold-colored, sans-serif font with a slight 3D effect and shadow.

what does increase performance means?

- for given cost, running an instance on arm can save you more by increasing throughput.

why run db on arm?



what does increase performance means?

- for given cost, running an instance on arm can save you more by increasing throughput.

cost of instance: 800 usd (yearly) [mysql-8.0.21 uniform/cpu-bound]

threads	arm (tps)	tps/usd (arm)	x86 (tps)	(tps/usd) (x86)	%
1024 (update-index)	78869	98 tps/usd	55537	69 tps/usd	42%

why run db on arm?



what does increase performance means?


- for given cost, running an instance on arm can save you more by increasing throughput.

**More TPS
@ Same Cost**

cost of instance: 800 usd (yearly) [mysql-8.0.21 uniform/cpu-bound]


threads	arm (tps)	tps/usd (arm)	x86 (tps)	(tps/usd) (x86)	%
1024 (update-index)	78869	98 tps/usd	55537	69 tps/usd	42%

why run db on arm?




“Can I scale DB w/o migrating database and keeping the cost constant?”

why run db on arm?



“Can I scale DB w/o migrating database and keeping the cost constant?”

Need to scale MySQL?
Go get new, better, larger
hardware/vm
@ EXTRA COST




why run db on arm?



“Can I scale DB w/o migrating database and keeping the cost constant?”


- #mysqlonarm

Need to scale MySQL?
Go get new, better, larger
hardware/vm
@ EXTRA COST




Switch to ARM.
Same Cost, Same DB,
Same environment.
Just More Performance.

why run db on arm?



“More throughput and also helps mother earth”


Need to scale MySQL?
Go get new, better, larger
hardware/vm
@ EXTRA COST




Switch to ARM.
Same Cost, Same DB,
Same environment.
👉 + More Performance.



why run db on arm?

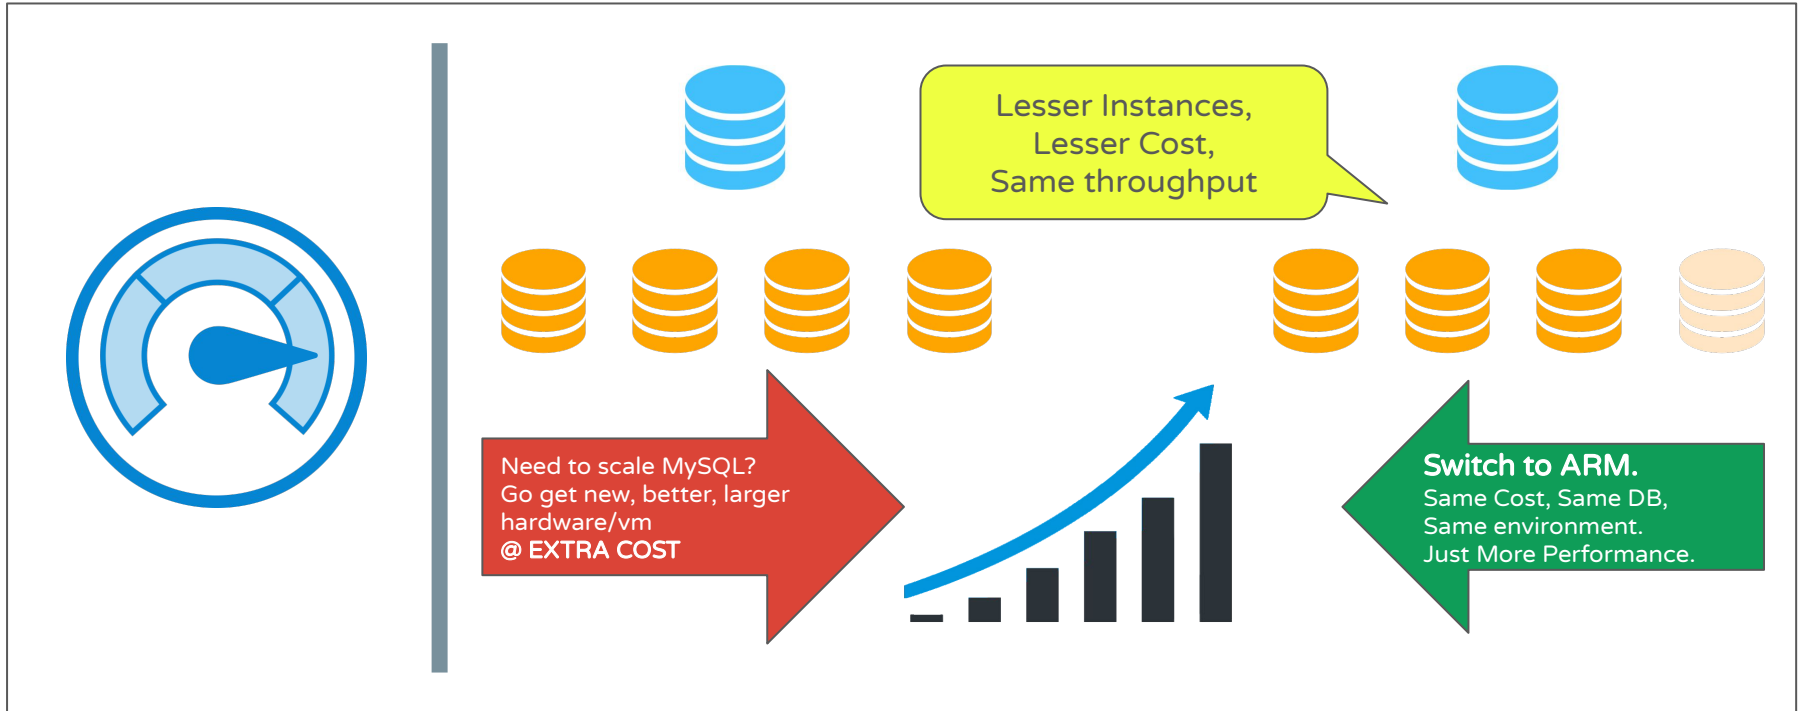


Need to scale MySQL?
Go get new, better, larger
hardware/vm
@ EXTRA COST




Switch to ARM.
Same Cost, Same DB,
Same environment.
Just More Performance.

why run db on arm?



why run db on arm?



- Any flavor (mysql/percona/mariadb)
- Any workload (cpu/io bound)
- Any pattern (read-only/read-write)
- Any configuration (contention/non-contention)
- Any setup (single node/multi-node)

ARM always score

agenda

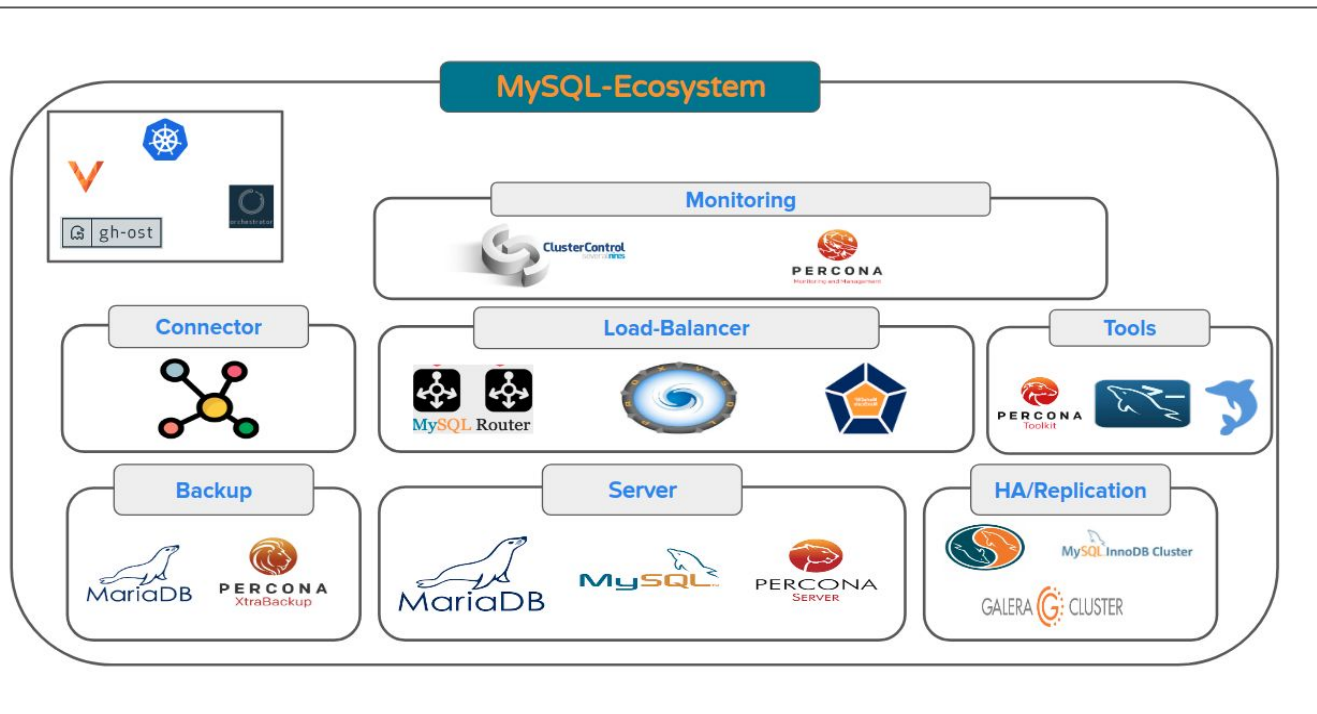
- growing arm ecosystem
- mysql/mariadb/percona on arm
- why run dbs on arm?
- state of ecosystem
- moving forward

state of ecosystem



- all that is good but what about ecosystem/supporting tools/support/community?
 - mysql/mariadb officially support server so rest assured your correctness issues (if any) would be looked on with same priority like x86.
 - db on arm in general is gaining lot of traction and mysql ecosystem too is adding users/developers expanding #arm community.

state of ecosystem



state of ecosystem



- server: **mysql** **mdb** **ps**
- load-balancer: **proxysql** **maxscale** **mysql-router**
- monitoring: **pmm** **cluster-control**
- ha/replication: **binlog**/**galera-cluster**/**group replication**
- backup: **mariabackup** **percona-backup**
- connector: **mysql-connector** **mariadb-connector**
- toolkit: **mysql-shell** **percona-toolkit**

for some tools community is trying build tools on arm.
with more support and push from community we can
expect ecosystem to build faster.

“Community need your support”



state of ecosystem









- why aim for completion?
 - server can run on arm and tools on x86 to start with but all things on arm can help realize better saving
 - also, we are hoping other environmental components like sharding solution, migration tools, containers can join the ecosystem.

agenda

- growing arm ecosystem
- mysql/mariadb/percona on arm
- why run dbs on arm?
- state of ecosystem
- moving forward

let's remain connected

- read more about **#mysqlonarm**
 - <https://mysqlonarm.github.io/>
- slack channel
 - mysql-slack: **#mysqlonarm** 
 - mariadb-zulip-chat: **#mariadbbonarm** 
- reach me
 -  krunalbauskar@gmail.com
 -  mysqlonarm@gmail.com
- social
 -  [#mysqlonarm](https://twitter.com/mysqlonarm)
 -  <https://in.linkedin.com/in/krunal-bauskar-b7a0b66>

q&a

- so when are you moving

