MariaDB Backup Compression with QAT

XPS SW EMEA

Marcin Olszewski

MariaDB Backup Compression with QAT

Testing impact of QAT compression on MariaDB backup

- Running compression under write-heavy load on MariaDB
- Full backup performed on all databases
- HammerDB TPROC-C benchmark generates write load on MariaDB
- Backup is performed with mariabackup producing xbstream, compressed externally with either parallel gzip or QAT gzip utility (qzip)

Comparison of CPU Load with QZIP vs Gzip

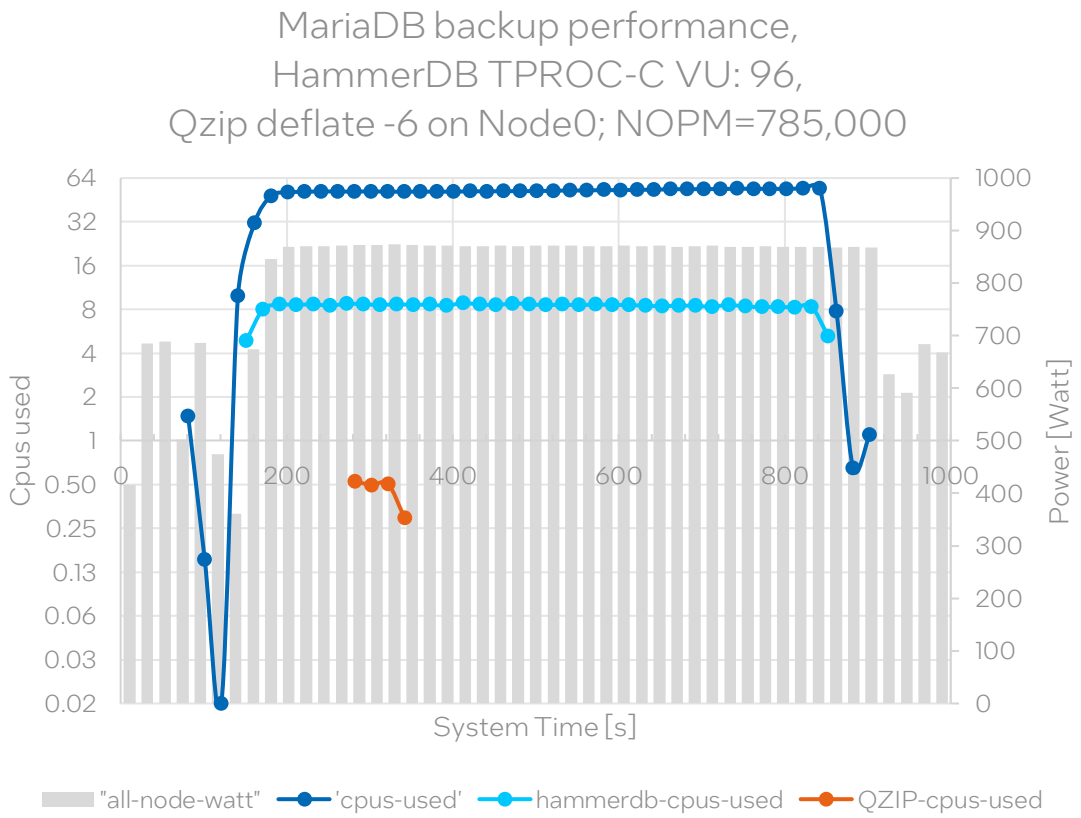

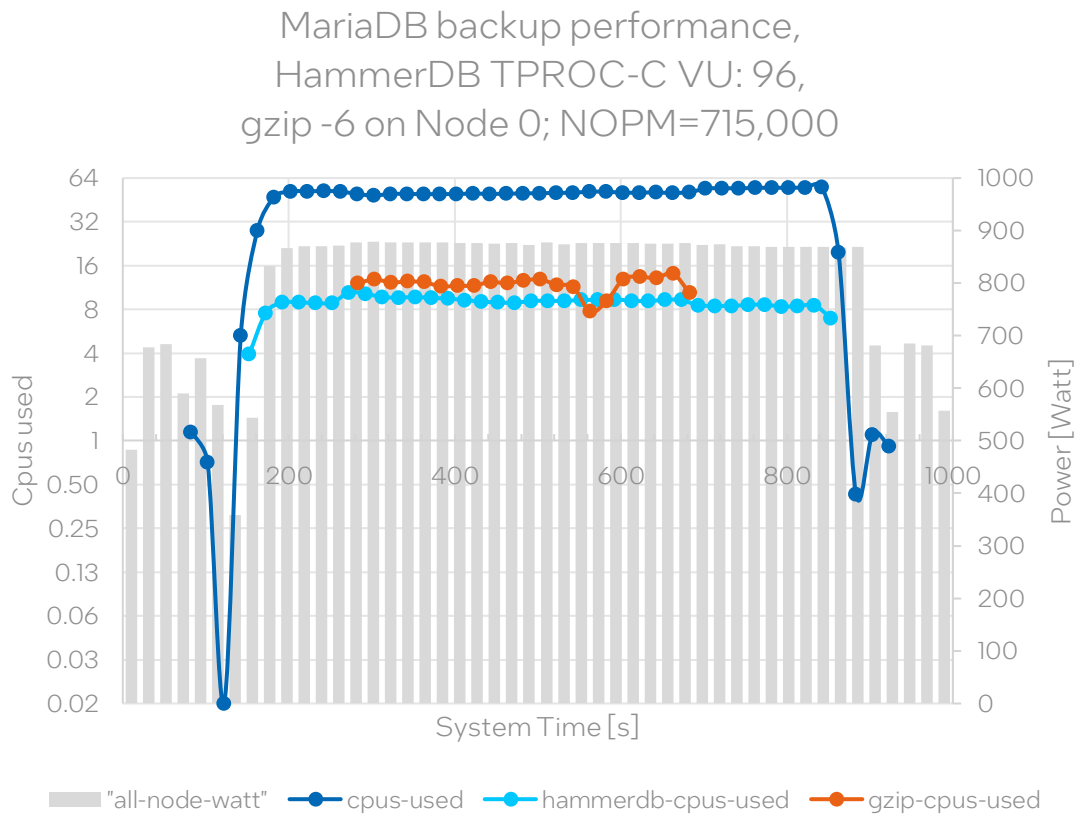

Performance of DB under backup with QZIP vs Gzip

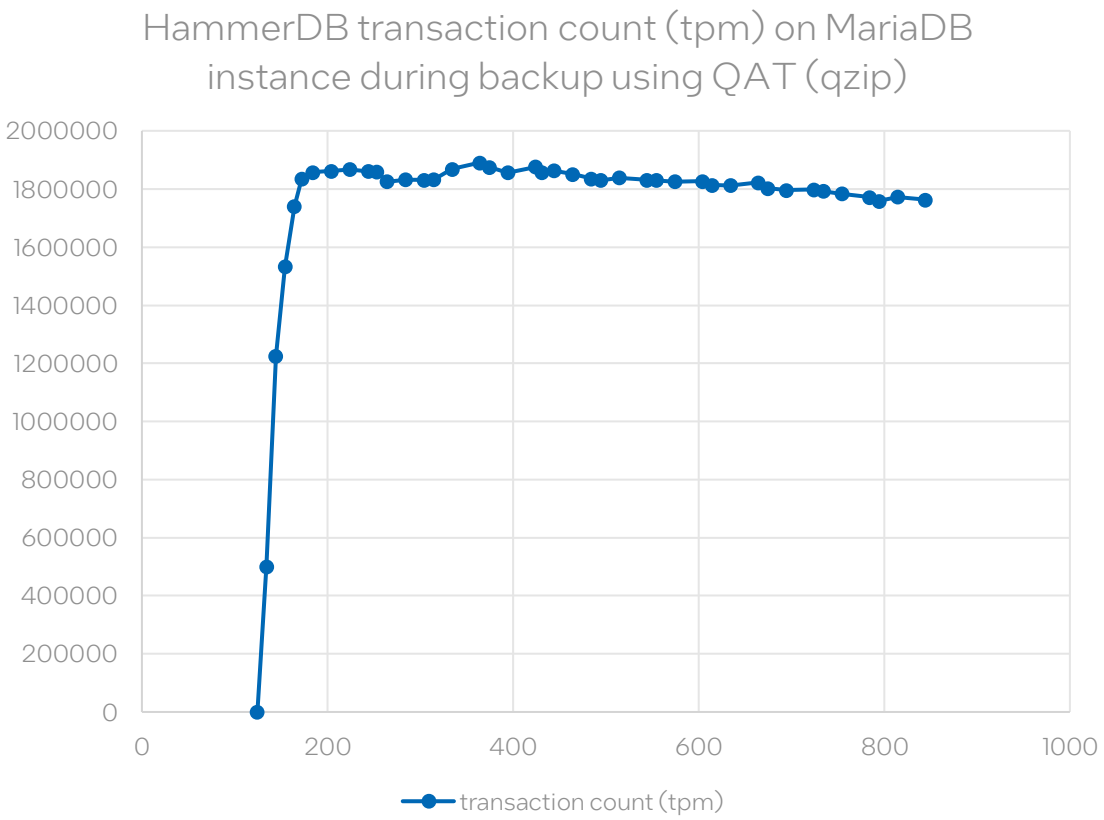

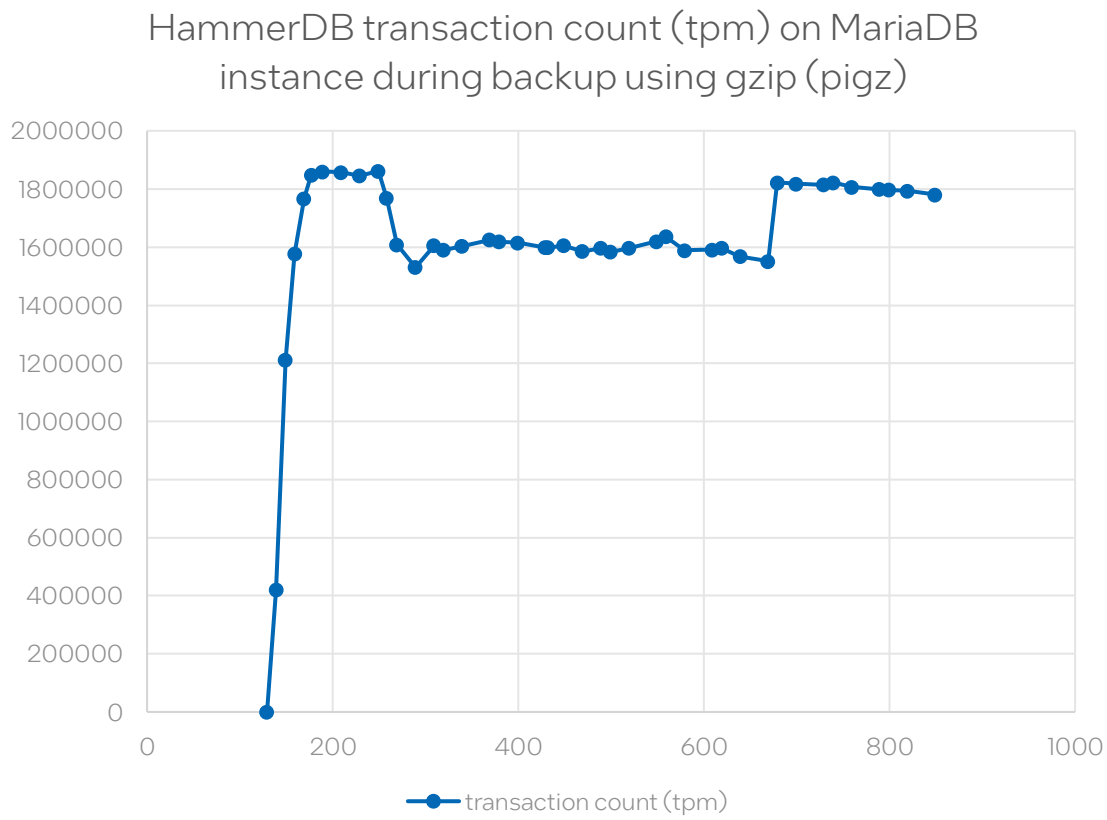

Most notable characteristics

QAT compression

- Backup time: 80s
- Average CPU utilization: 0.4
- 99% of baseline NOPM during backup
- offloading compression to QAT releases all CPU time for DB

GZIP compression

- Backup time: 420s
- Average CPU utilization: 12.0
- 90% of baseline NOPM during backup

System Configuration

	Config 1-2
Name	001b21e7c6ba.jf.intel.com
Time	Wed Dec 13 12:27:01 PM UTC 2023
System	Intel Corporation ArcherCityM
Baseboard	Intel Corporation ArcherCityM
Chassis	..... Rack Mount Chassis
CPU Model	Intel(R) Xeon(R) Platinum 8468H
Microarchitecture	SPR_XCC
Sockets	4
Cores per Socket	48
Hyperthreading	Enabled
CPUs	384
Intel Turbo Boost	Enabled
Base Frequency	2.1GHz
All-core Maximum Frequency	3.0GHz
Maximum Frequency	3.8GHz
NUMA Nodes	4
Prefetchers	L2 HW, L2 Adj., DCU HW, DCU IP
PPINs	561de81f89e69656,560ef71f655ed053,561de91f97fdef96,560efe1ff30835f2
Accelerators (4th Gen Only)	DLB 16 [0], DSA 16 [0], IAA 16 [0], QAT 16 [0]
Installed Memory	480GB (30x16GB DDR5 4800 MT/s [4800 MT/s])
Hugepagesize	2048 kB
Transparent Huge Pages	madvise
Automatic NUMA Balancing	Enabled
NIC	2x I210 Gigabit Network Connection
Disk	1x 1.5T INTEL SSDSC2BB01, 1x 745.2G INTEL SSDSC2BA80
BIOS	EGSDCRB1.SYS.0084.D24.2207132145
Microcode	0x2b0004d0
OS	Ubuntu 22.04.3 LTS
Kernel	5.15.0-89-generic
TDP	330 watts
Power & Perf Policy	Normal (6)
Frequency Governor	powersave
Frequency Driver	intel_pstate
Max C-State	9

Software/Workload Configuration

Category	Attribute	Config1	Config2
Run Info			
	Benchmark		
	Test by	marcinol	marcinol
Software			
	Workload	HammerDB 4.8	HammerDB 4.8
	OS	Ubuntu 22.04.3 LTS	Ubuntu 22.04.3 LTS
	Kernel	5.15.0-89-generic	5.15.0-89-generic
Workload Specific Details			
	Workload	HammerDB 4.8 TPROC-C	HammerDB 4.8 TPROC-C
	Database	MariaDB 11.3	MariaDB 11.3
	Database Storage	180GiB RAM (tmpfs), bound outside DB NUMA Node	180GiB RAM FS (tmpfs), bound outside DB NUMA Node
	Run Method	2 min warmup, 10 min TPROC-C	2 min warmup, 10 min TPROC-C
	Compression Method	QZIP deflate level 6	GZIP deflate level 6
	Workload Configuration	400 Warehouses 96 Virtual Users	400 Warehouses 96 Virtual Users
	Command Line	mariabackup -u sysbackup -S /tmp/mysql.sock --backup --stream=xbstream --no-lock | pv --quiet --rate-limit=1024M | qzip -L6 -O gzip -f >db-backup.gz	mariabackup -u sysbackup -S /tmp/mysql.sock --backup --stream=xbstream --no-lock | pv --quiet --rate-limit=1024M | pigz -c -6 >db-backup.gz

Benchmark Setup

- Deployed MariaDB off repository branch 11.3
- HammerDB version 4.8
- Backup with compression command line:
  - (gzip)
```
mariabackup --backup --stream=xbstream --no-lock | pv --rate-limit=1024M \
| pigz -c -6
```
  - (QAT)
```
mariabackup --backup --stream=xbstream --no-lock | pv --rate-limit=1024M \
| qzip -L6 -O gzip -f
```
  - (all programs bound to numa-socket 0)
- To create comparable conditions for backup with QAT and GZIP used pv for rate-limiting transfer between backup and compression utility
- GZIP was unable to keep up with transfer, thus the backup time with GZIP was much longer (420 vs 80s)
- Backup process was only started after the workload warmup phase (2 minutes) was over


intel