

8.9.3 against 8.9.1-rerun  
value-only workload

Michael Karg, Cardano Performance team

2024-05-23

# Contents

<b>1</b>	<b>Manifest</b>	<b>2</b>
<b>2</b>	<b>Analysis</b>	<b>4</b>
2.1	Resource Usage . . . . .	4
2.2	Anomaly control . . . . .	4
2.3	Forging . . . . .	5
2.4	Individual peer propagation . . . . .	5
2.5	End-to-end propagation . . . . .	5
<b>I</b>	<b>Appendix A: charts</b>	<b>6</b>
<b>3</b>	<b>Cluster performance charts</b>	<b>7</b>
<b>II</b>	<b>Appendix B: data dictionary</b>	<b>25</b>
<b>4</b>	<b>Block propagation metrics</b>	<b>26</b>
<b>5</b>	<b>Cluster performance metrics</b>	<b>28</b>

# Chapter 1

## Manifest

We compare 8.9.3 (Babbage) relative to 8.9.1-rerun (Babbage), under value-only workload.

	8.9.1-rerun	8.9.3
Analysis date	2024-05-21	2024-05-09
Cluster system start date	2024-05-20	2024-05-08
Cluster system start time	07:55:46	13:39:09
Identifier	8.9.1	8.9.3
Run batch	891rerun	8.9.3
GHC version	8.10.7	8.10.7
cardano-node version	8.9.1	8.9.3
ouroboros-consensus version	0.16.0.0	0.16.0.0
ouroboros-network version	0.13.1.0	0.15.0.0
cardano-ledger-core version	1.10.0.0	1.10.0.0
plutus-core version	1.21.0.0	1.21.0.0
cardano-crypto version	1.1.2	1.1.2
cardano-prelude version	0.1.0.4	0.1.0.4
cardano-node git	07524f8	be18440
ouroboros-consensus git	a2cb6e5	a2cb6e5
ouroboros-network git	e8d1721	6a947bc
cardano-ledger-core git	6e2d37c	6e2d37c
plutus-core git	022595e	022595e
cardano-crypto git	6568a5e	6568a5e
cardano-prelude git	a6f18f7	a6f18f7
Era	babbage	babbage
Delegation map size	1000000	1000000
Starting UTxO set size	4000000	4000000
Extra tx payload	100	100
Tx inputs	2	2
Tx Outputs	2	2
TPS	12.0	12.0
Transaction count	768000	768000
Plutus script	—	—
Machines	52	52
Number of filters applied	3	3
Log text lines emitted per host	5150057.8461	5568781.8653
Log objects emitted per host	5150027.8461	5568751.8653
Log objects analysed per host	2284205.0	2436307.6538
Host run time, s	63905.7	63928.9
Host log line rate, Hz	80.588	87.109
Total log objects analysed	118778660	126687998
Run time, s	63912	63936
Analysed run duration, s	48016	48035
Run time efficiency	0.75	0.75
Node start spread, s	10.277754	11.509614
Node stop spread, s	4.3709770	4.2418675
Perf analysis start spread, s	0	0
Perf analysis stop spread, s	4	5
Slots analysed	48013	48031
Blocks analysed	2242	2194
Blocks rejected	871	952

# Chapter 2

## Analysis

### 2.1 Resource Usage

	8.9.1-rerun	8.9.3	$\Delta$	$\Delta\%$
Forge loop starts, #	0.99872	0.99872	0.000	0
Process CPU usage, %	8.2654	8.5797	0.314	4
RTS GC CPU usage, %	1.1615	1.2491	0.088	8
RTS Mutator CPU usage, %	7.0978	7.3215	0.224	3
Major GCs, #	0.00099	0.00099	0.000	0
Minor GCs, #	2.1095	2.1728	0.063	3
Kernel RSS, MB	8297.8	8307.7	9.900	0
RTS heap size, MB	8246.8	8256.6	9.800	0
RTS live GC dataset, MB	3797.5	3796.5	-1.000	0
RTS alloc rate, MB/s	65.895	67.885	1.990	3
Filesystem reads, KB/s	9e-05	0.00124	0.001	1111
Filesystem writes, KB/s	231.79	228.27	-3.520	-2
CPU 85% spans, slots	0.06456	0.07892	0.014	22
Sample count	(249>)	(249>)		

### 2.2 Anomaly control

	8.9.1-rerun	8.9.3	$\Delta$	$\Delta\%$
Blocks per host, blocks	61.923	62.692	0.769	1
Filtered to chained block ratio, /	0.7216	0.69763	-0.024	-3
Chained to forged block ratio, /	0.96636	0.96476	-0.002	0
Height & slot battles, blocks	0.00223	0.00501	0.003	135
Block size, B	89019	89004	-15	0
Sample count	(52)	(52)		

## 2.3 Forging

	8.9.1-rerun	8.9.3	$\Delta$	$\Delta\%$
Started forge loop iteration, s	0.00122	0.00171	0.000	0
Acquired block context, s	0.02282	0.0236	0.001	4
Acquired ledger state, s	7e-05	6e-05	-0.000	0
Acquired ledger view, s	3e-05	2e-05	-0.000	0
Leadership check duration, s	0.00043	0.00045	0.000	0
Ledger ticking, s	0.02159	0.02117	-0.000	0
Mempool snapshotting, s	0.06828	0.06798	-0.000	0
Leadership to forged, s	0.00142	0.00141	-0.000	0
Forged to announced, s	0.0007	0.00072	0.000	0
Forged to sending, s	0.00521	0.00577	0.001	19
Forged to self-adopted, s	0.07073	0.07169	0.001	1
Slot start to announced, s	0.11658	0.11716	0.001	1
Sample count	(2242)	(2194)		

## 2.4 Individual peer propagation

	8.9.1-rerun	8.9.3	$\Delta$	$\Delta\%$
First peer notice, s	0.11835	0.11909	0.001	1
First peer fetch, s	0.12749	0.12862	0.001	1
Notice to fetch request, s	0.00114	0.00123	0.000	0
Fetch duration, s	0.36759	0.36037	-0.007	-2
Fetches to announced, s	1e-05	0.0	-0.000	0
Fetches to sending, s	0.04133	0.04151	0.000	0
Fetches to adopted, s	0.07296	0.07386	0.001	1
Sample count	(2242)	(2194)		

## 2.5 End-to-end propagation

	8.9.1-rerun	8.9.3	$\Delta$	$\Delta\%$
0.50 adoption, s	0.65328	0.64816	-0.005	-1
0.80 adoption, s	1.0537	1.0464	-0.007	-1
0.90 adoption, s	1.0701	1.0678	-0.002	0
0.92 adoption, s	1.0735	1.0716	-0.002	0
0.94 adoption, s	1.0783	1.079	0.001	0
0.96 adoption, s	1.0849	1.0857	0.001	0
0.98 adoption, s	1.0945	1.0941	-0.000	0
1.00 adoption, s	1.1234	1.1181	-0.005	0
Sample count	(2242)	(2194)		

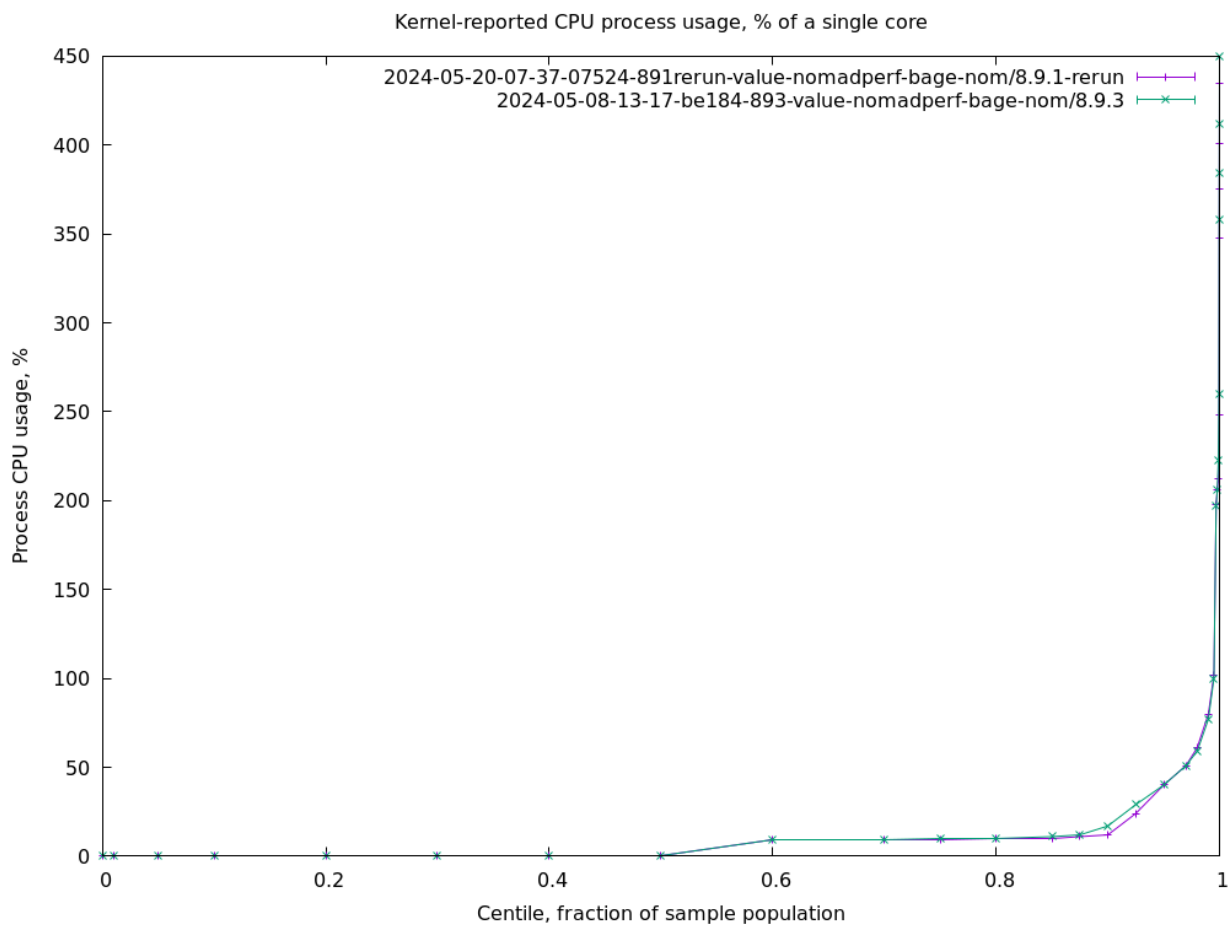
## Part I

# Appendix A: charts

# Chapter 3

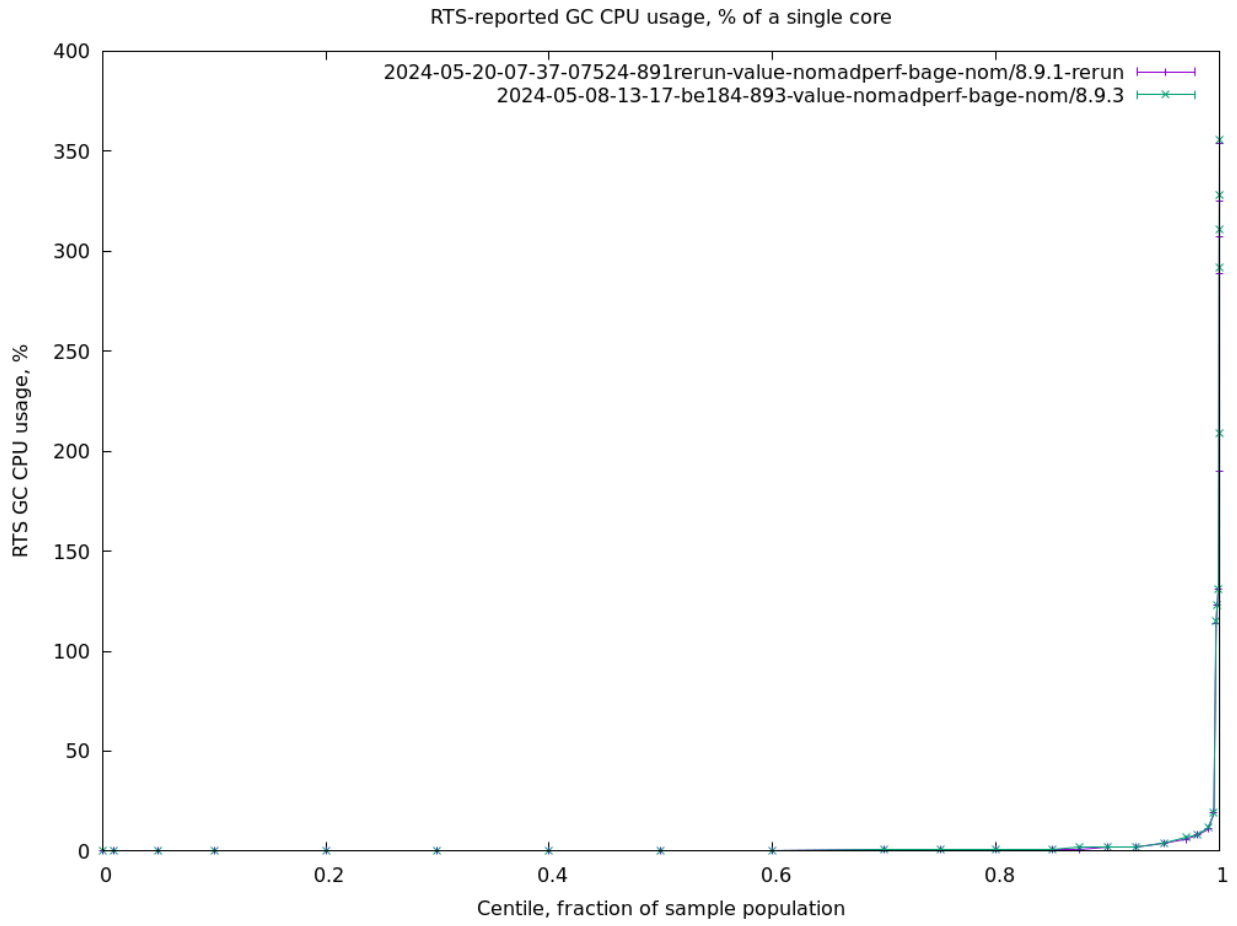
## Cluster performance charts

**Process CPU usage (CentiCpu)** Kernel-reported CPU process usage, % of a single core

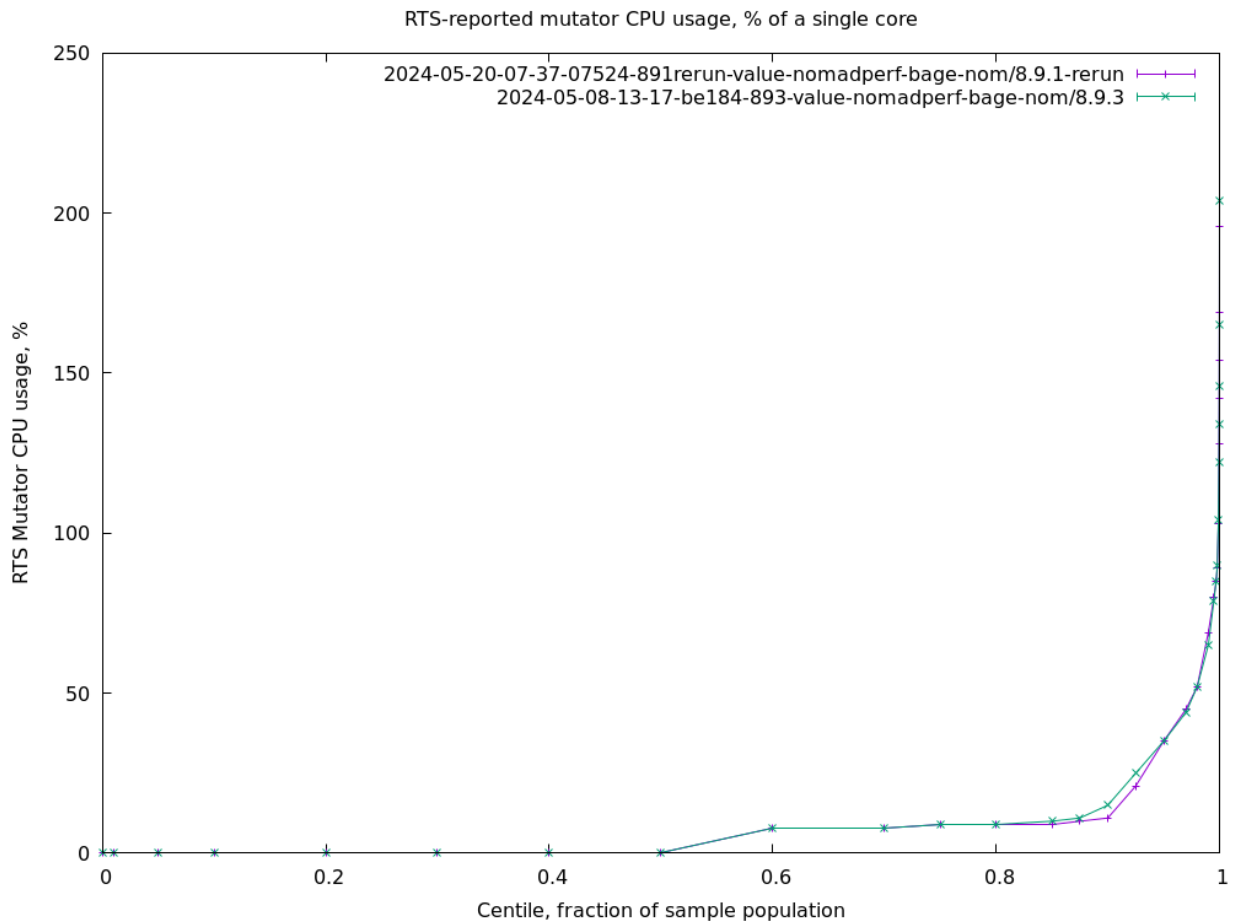


**RTS GC CPU usage (CentiGC)** RTS-reported GC CPU usage, % of a single core

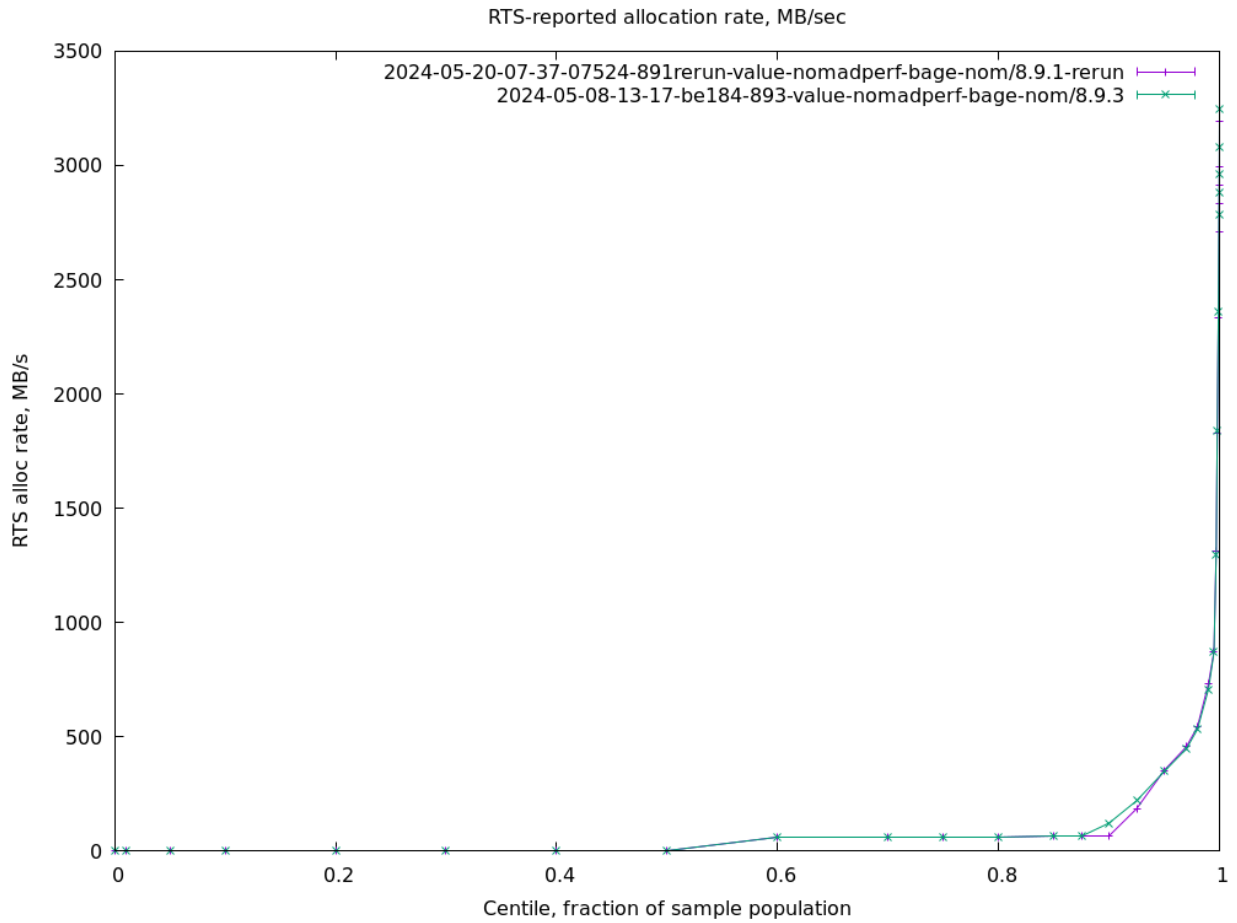




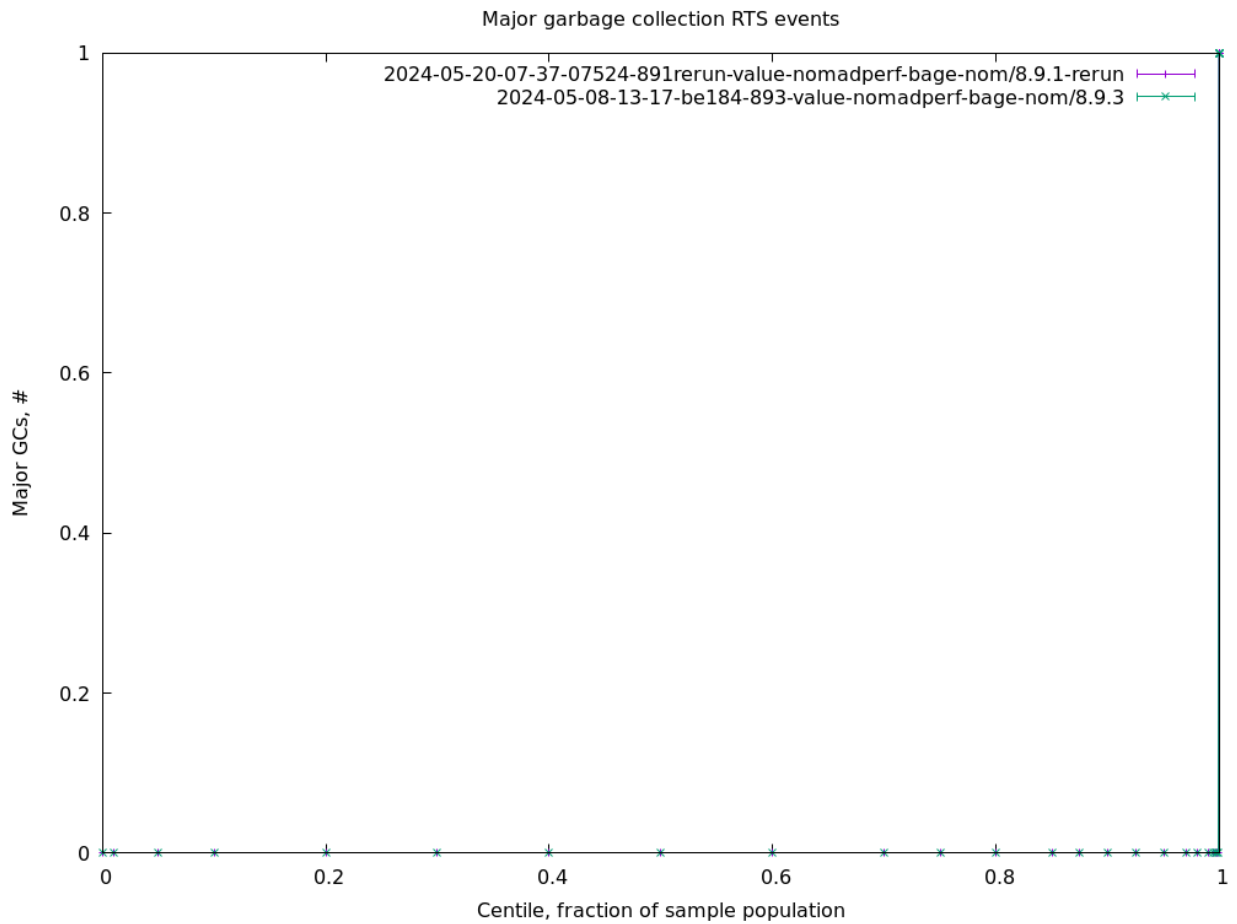
**RTS Mutator CPU usage (CentiMut)** RTS-reported mutator CPU usage, % of a single core



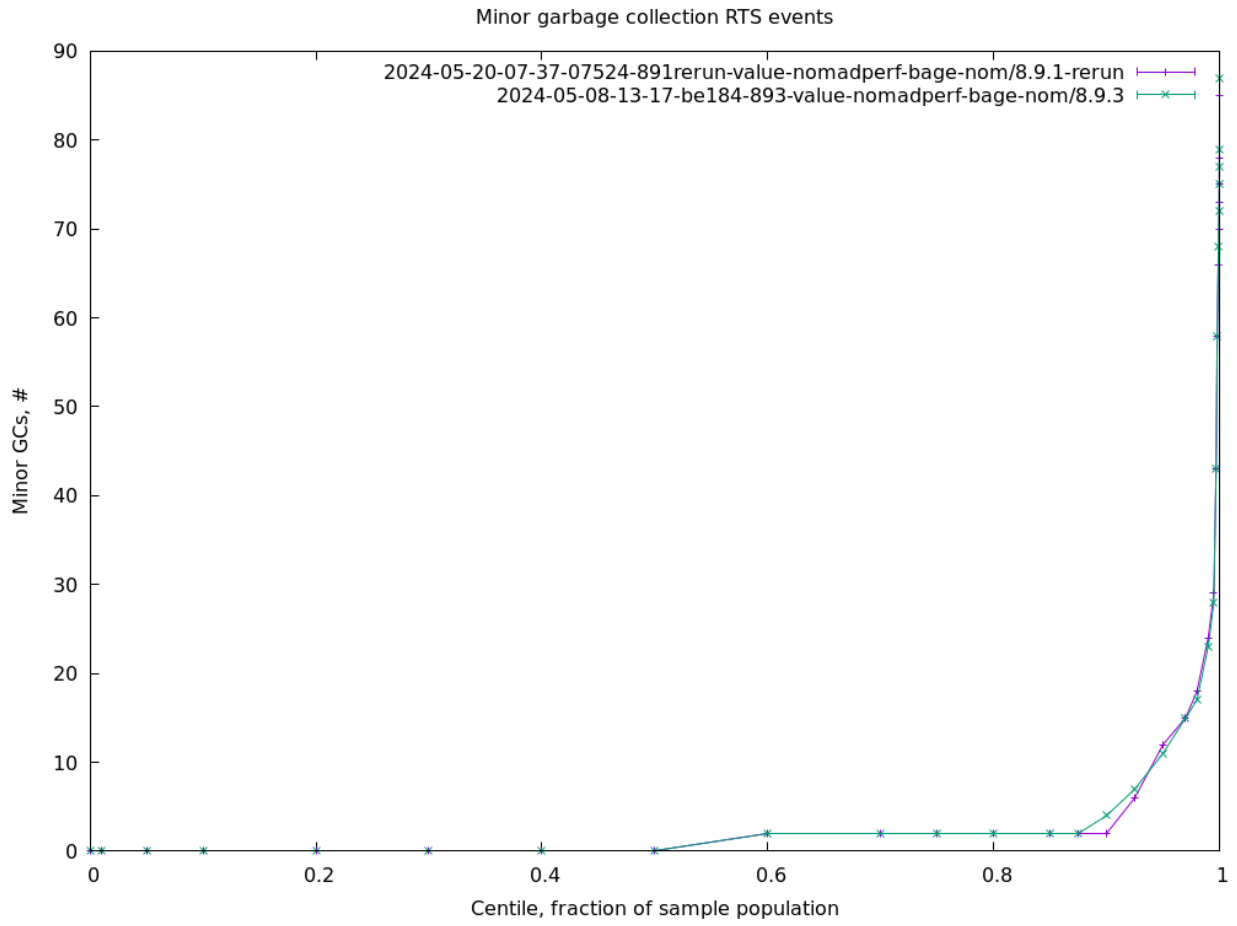
**RTS alloc rate (Alloc)** RTS-reported allocation rate, MB/sec



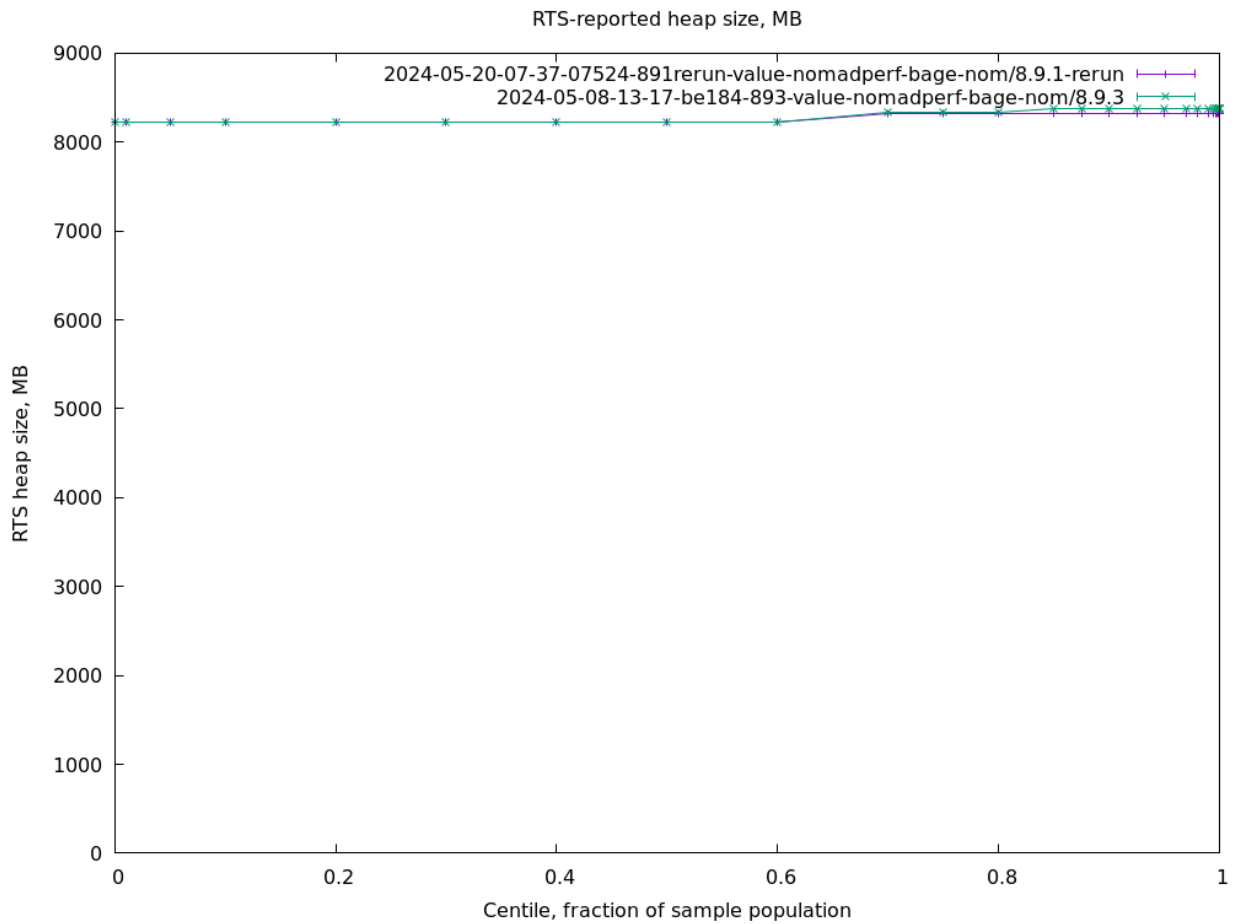
Major GCs (GcsMajor) Major garbage collection RTS events



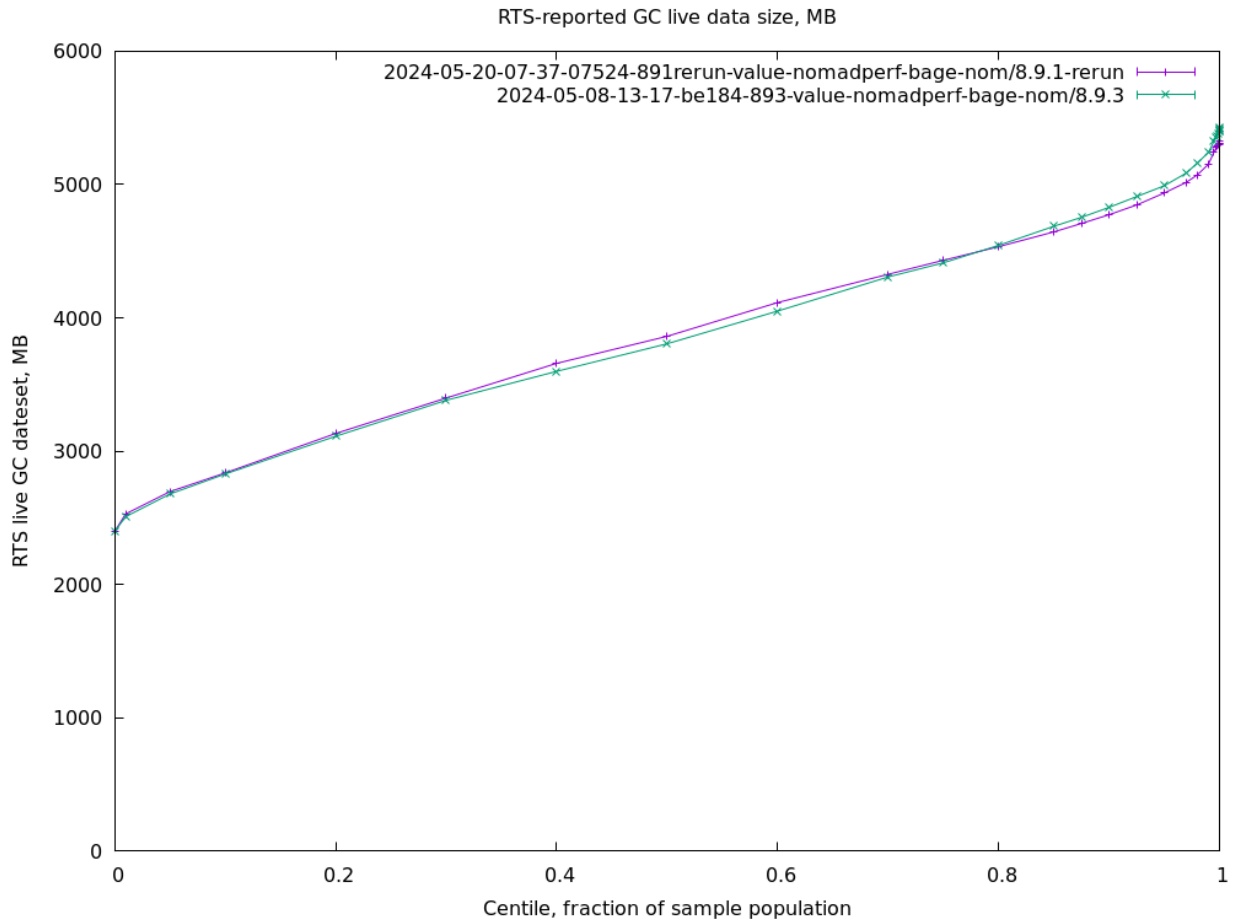
Minor GCs (GcsMinor) Minor garbage collection RTS events



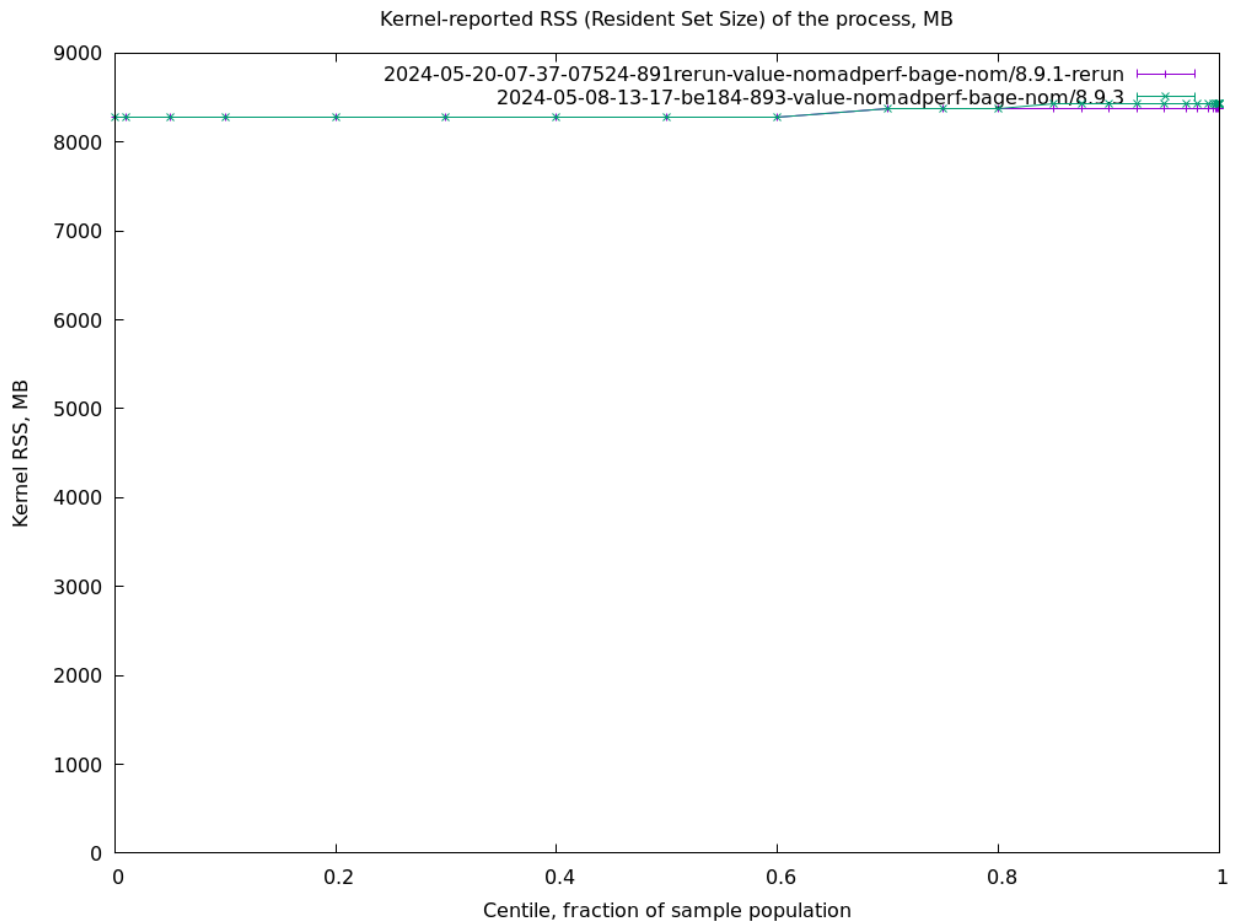
**RTS heap size (Heap)** RTS-reported heap size, MB



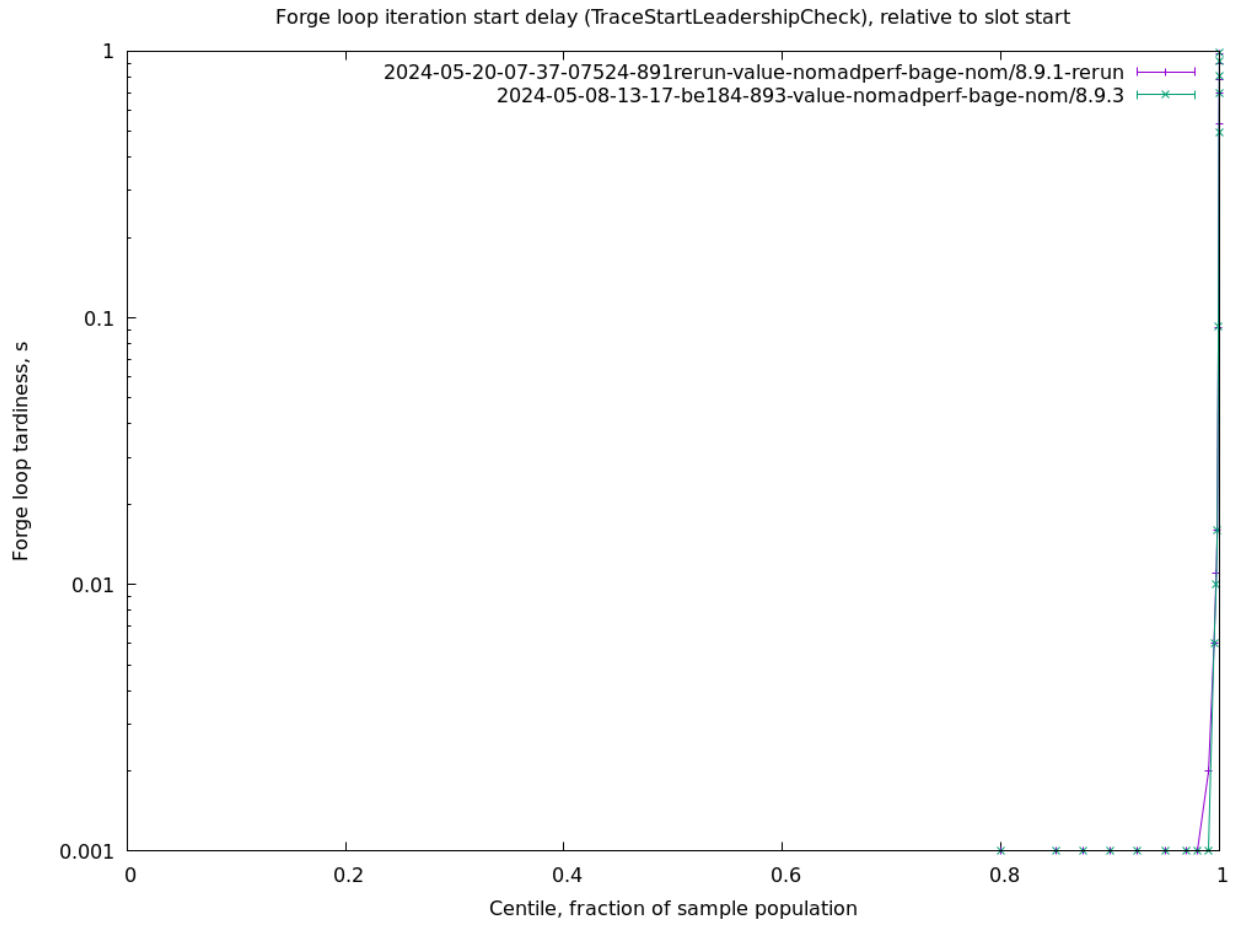
**RTS live GC dataset (Live)** RTS-reported GC live data size, MB



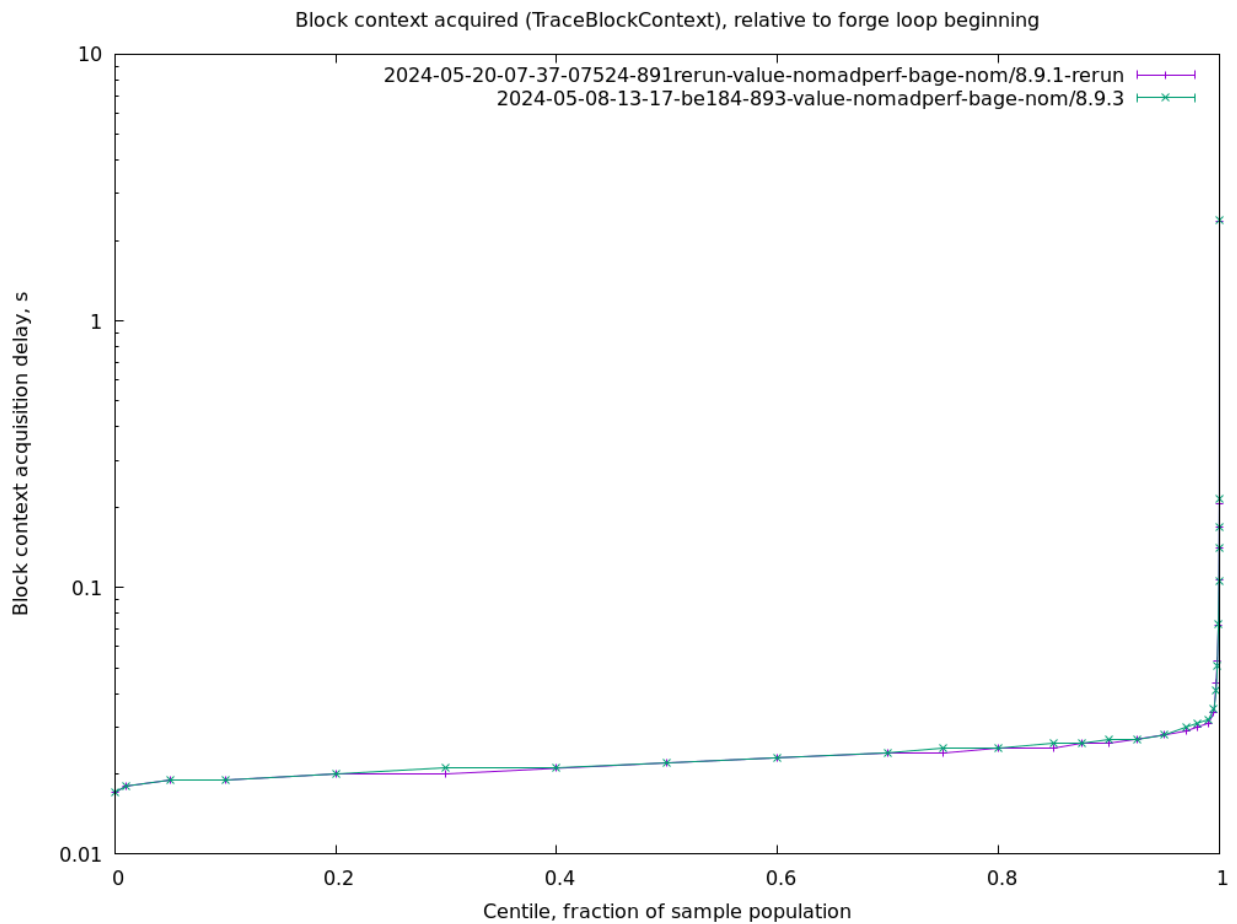
**Kernel RSS (RSS)** Kernel-reported RSS (Resident Set Size) of the process, MB



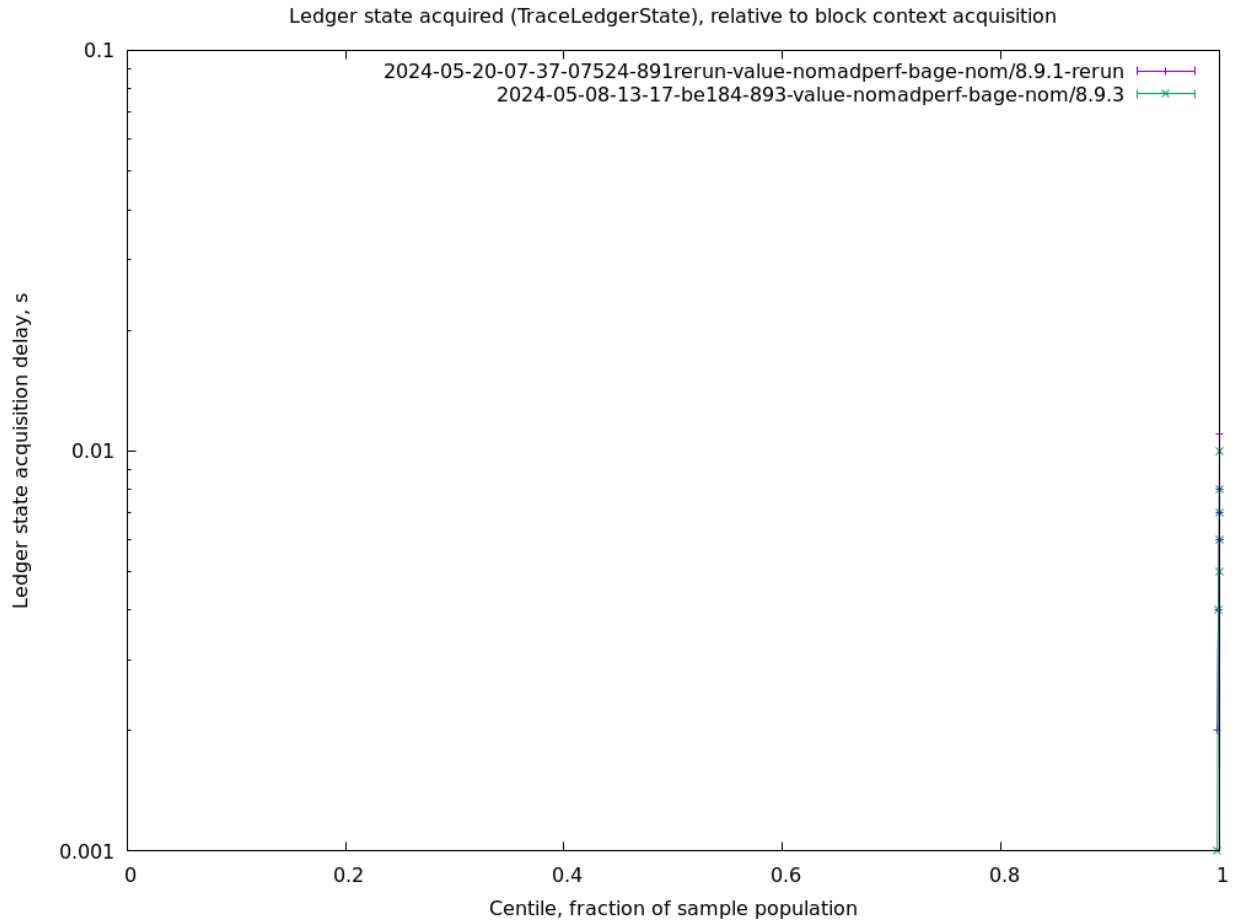
**Forge loop tardiness (cdfStarted)** Forge loop iteration start delay (TraceStartLeadershipCheck), relative to slot start



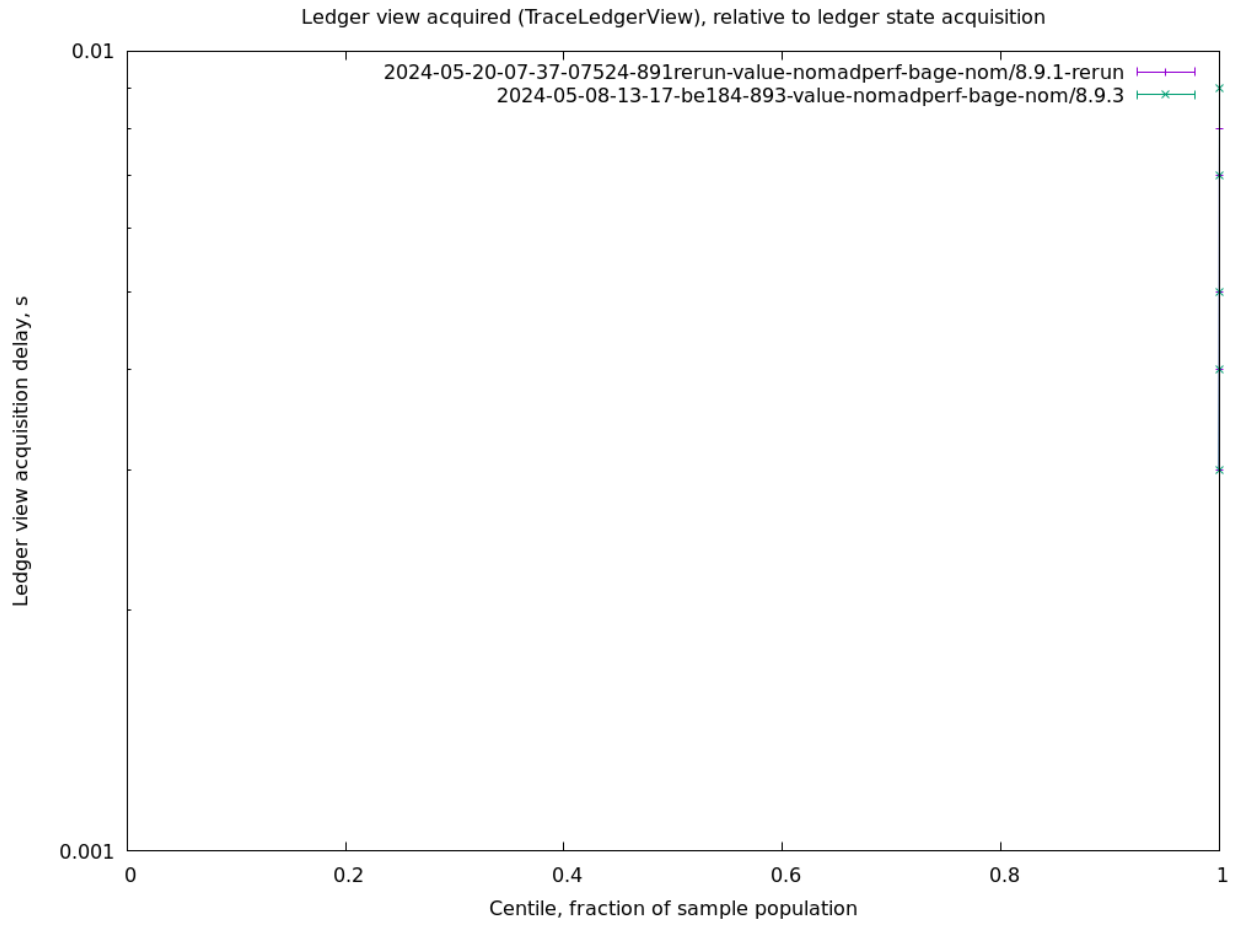
**Block context acquisition delay (cdfBlkCtx)** Block context acquired (TraceBlockContext), relative to forge loop beginning



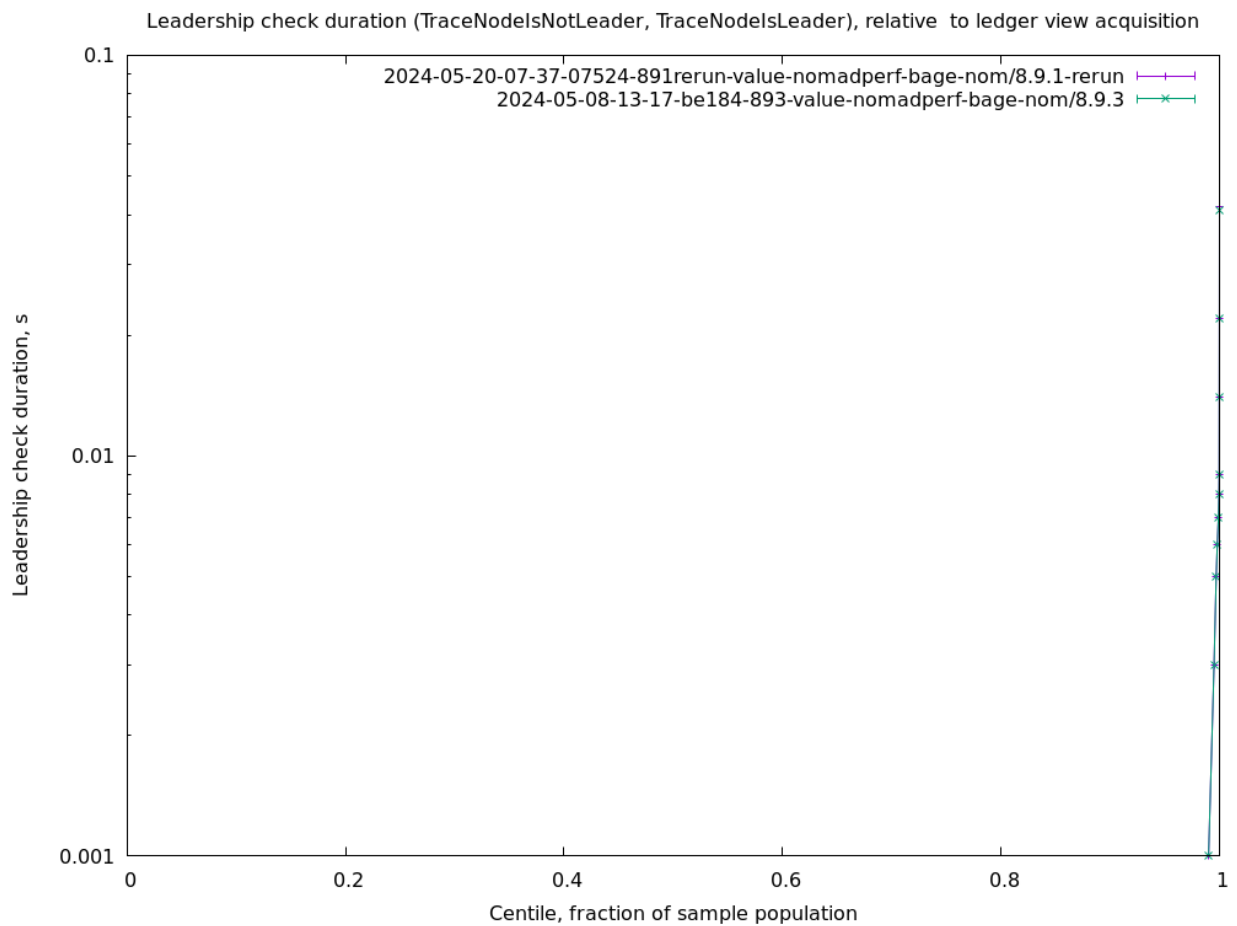
**Ledger state acquisition delay (cdfLgrState)** Ledger state acquired (TraceLedgerState), relative to block context acquisition



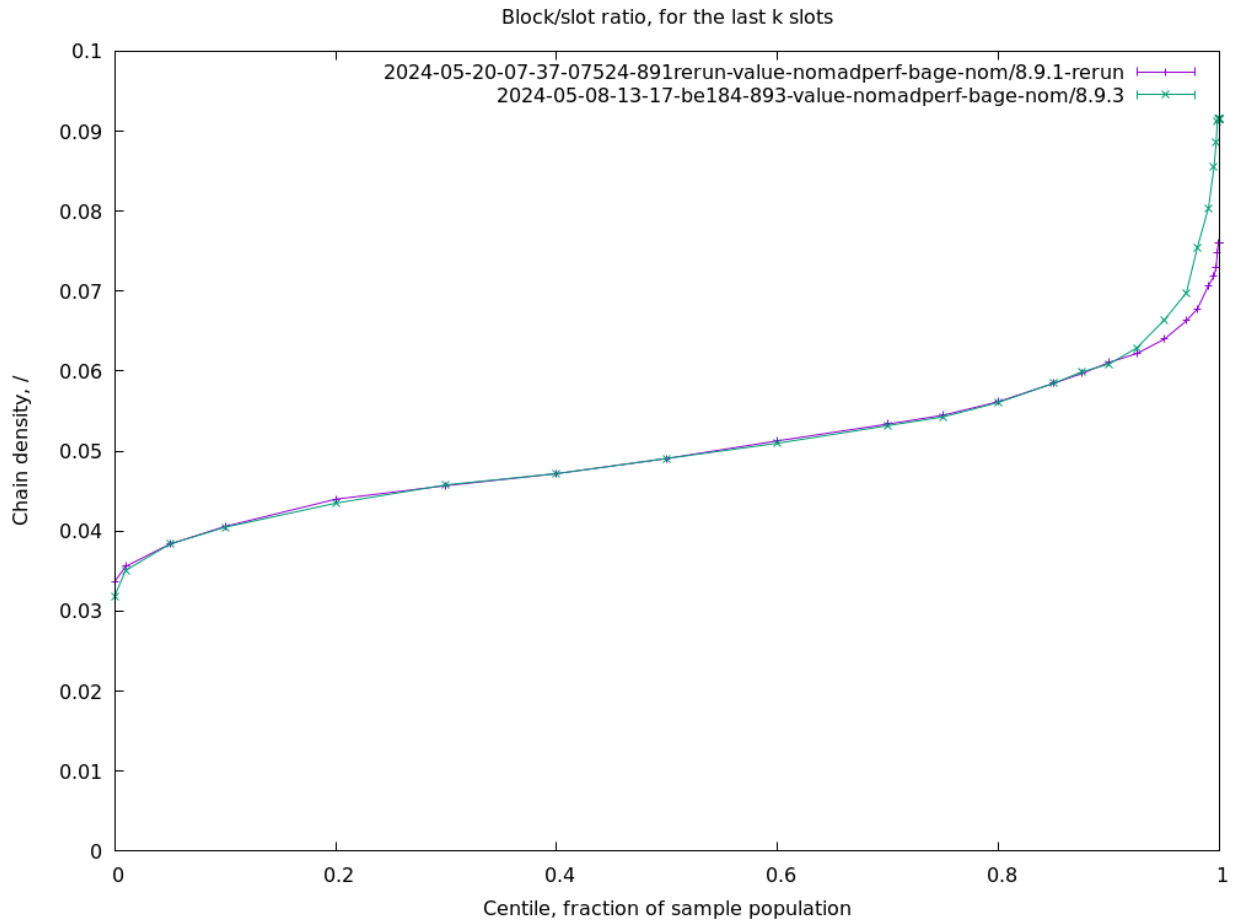
**Ledger view acquisition delay (cdfLgrView)** Ledger view acquired (TraceLedgerView), relative to ledger state acquisition



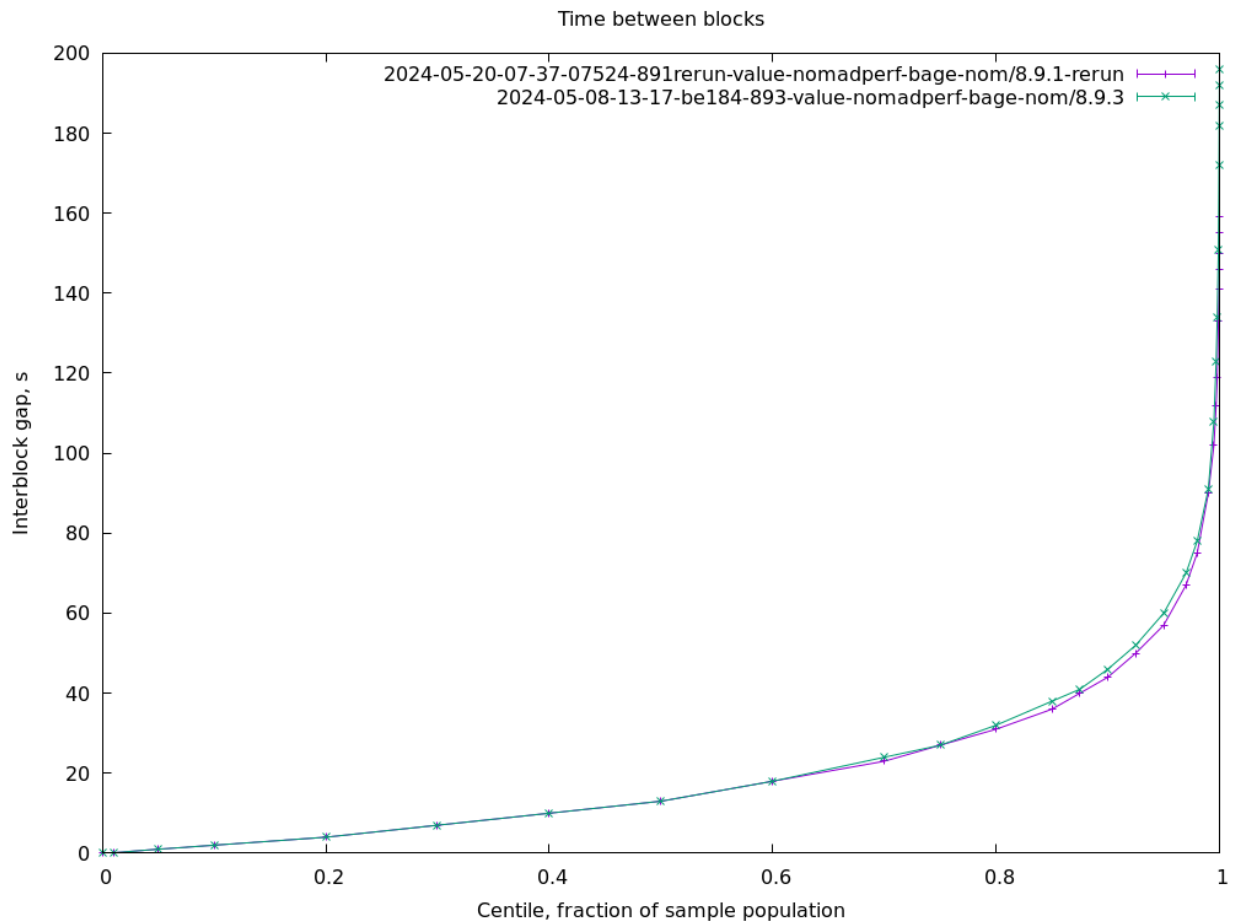
**Leadership check duration (cdfLeading)** Leadership check duration (TraceNodeIsNotLeader, TraceNodeIsLeader), relative to ledger view acquisition



**Chain density (cdfDensity)** Block/slot ratio, for the last 'k' slots

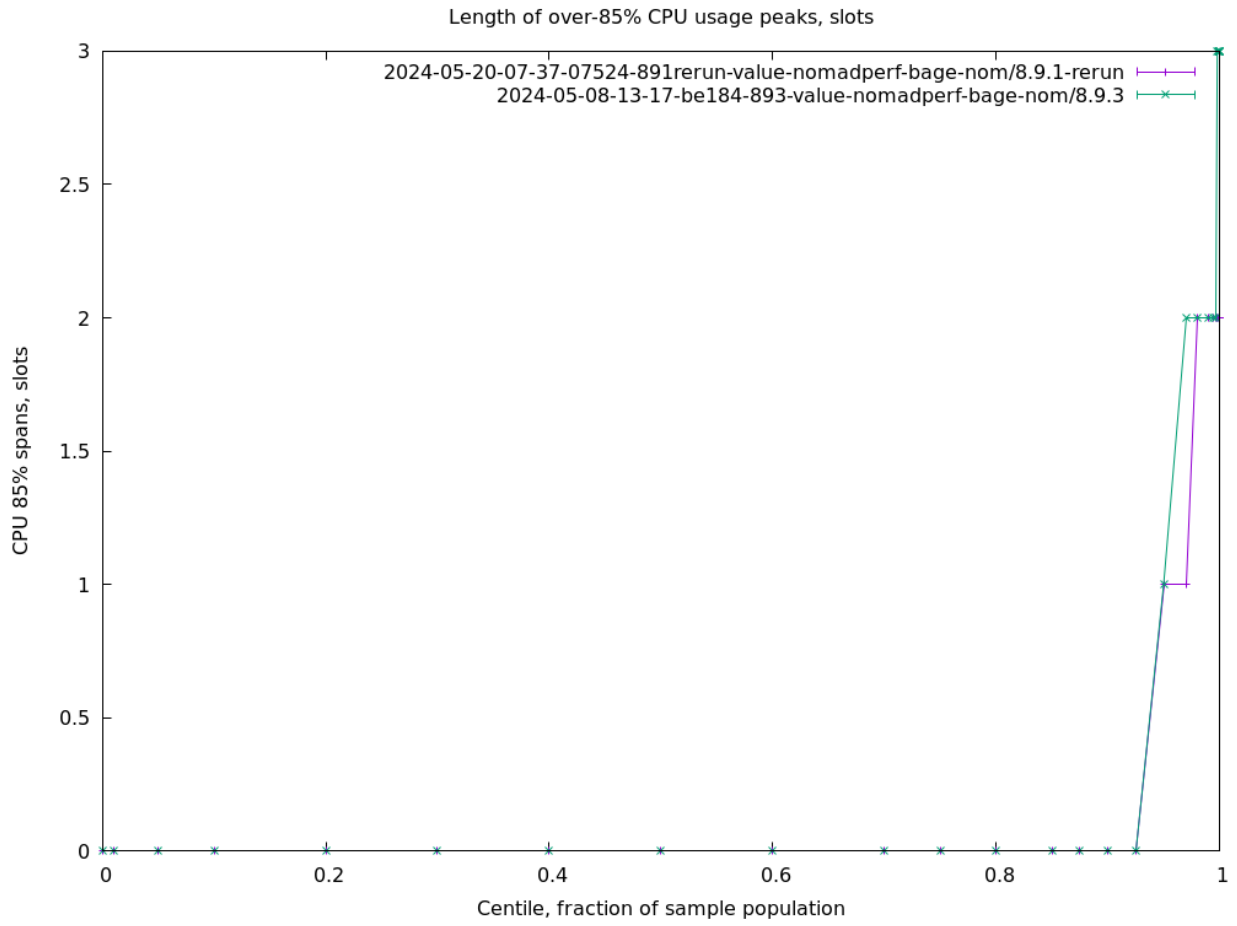


Interblock gap (cdfBlockGap) Time between blocks

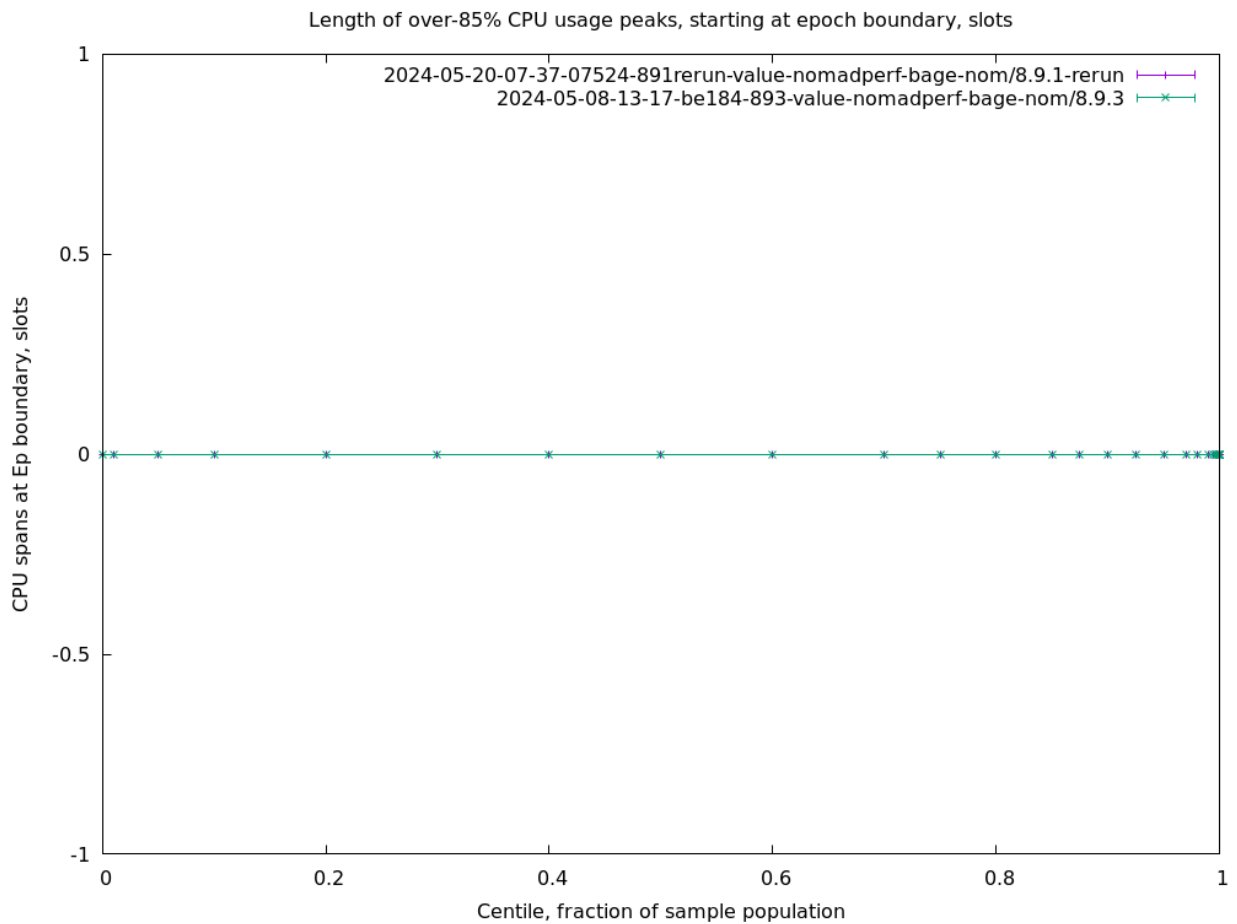


CPU 85% spans (cdfSpanLensCpu) Length of over-85% CPU usage peaks, slots

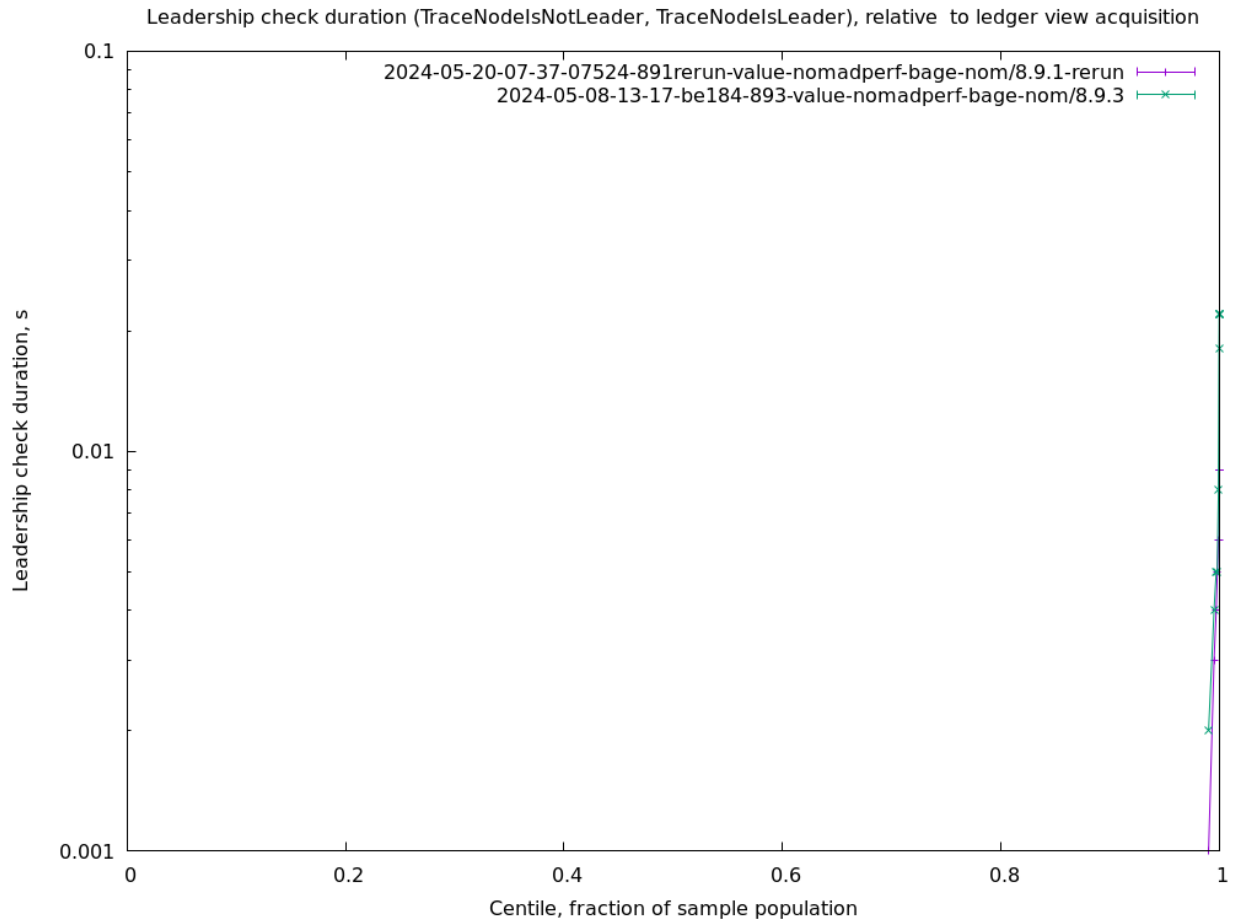




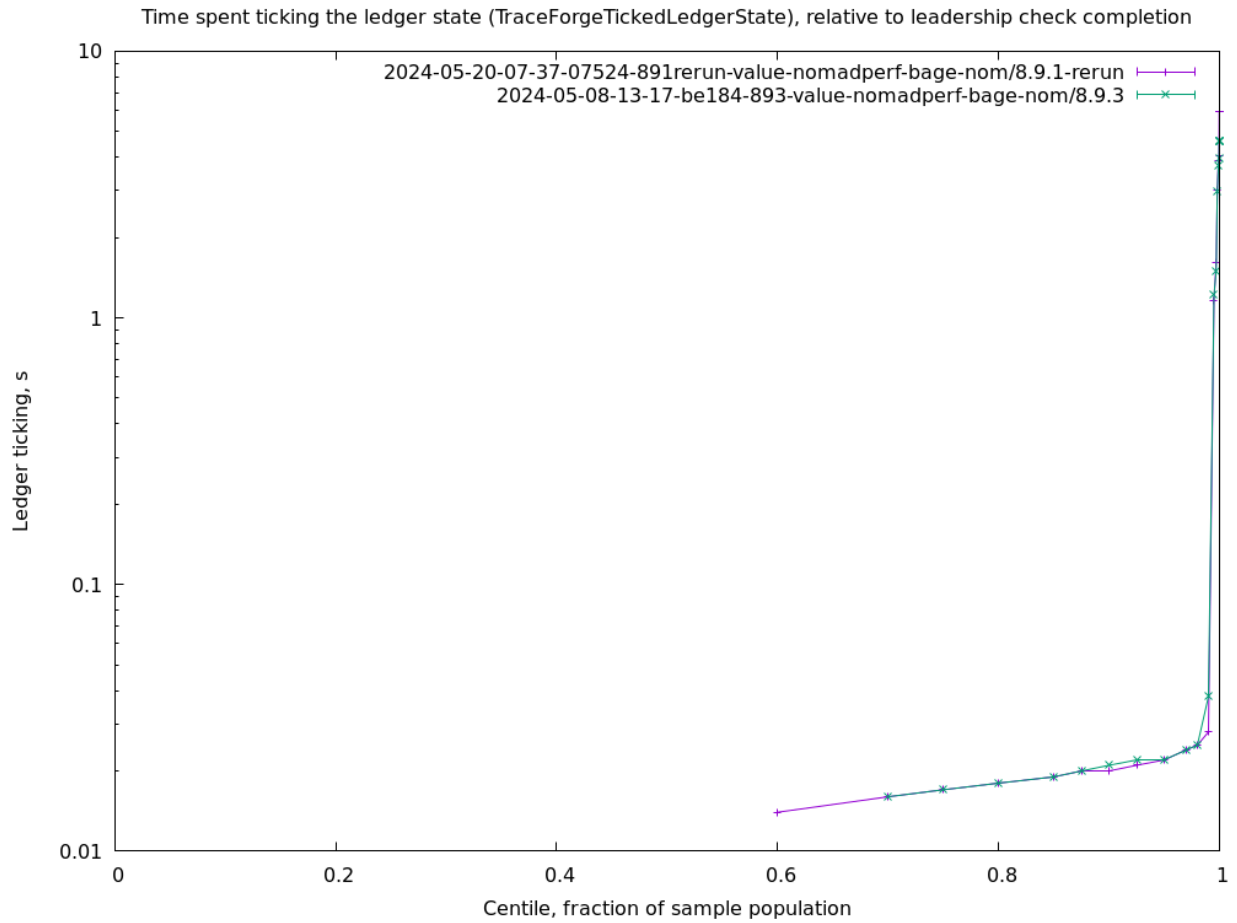
CPU spans at Ep boundary (cdfSpanLensCpuEpoch) Length of over-85% CPU usage peaks, starting at epoch boundary, slots



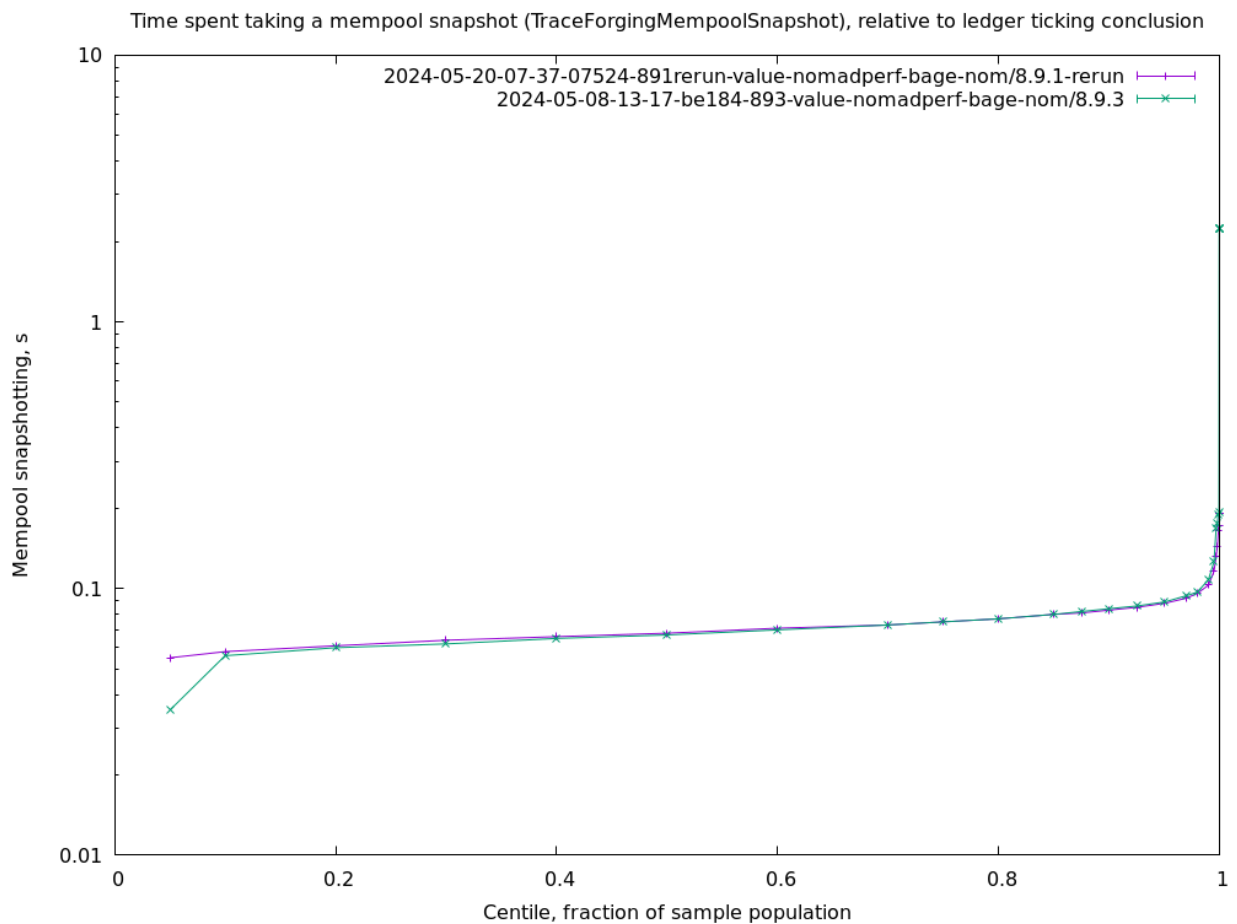
**Leadership check duration (cdfForgerLead)** Leadership check duration (TraceNodeIsNotLeader, TraceNodeIsLeader), relative to ledger view acquisition



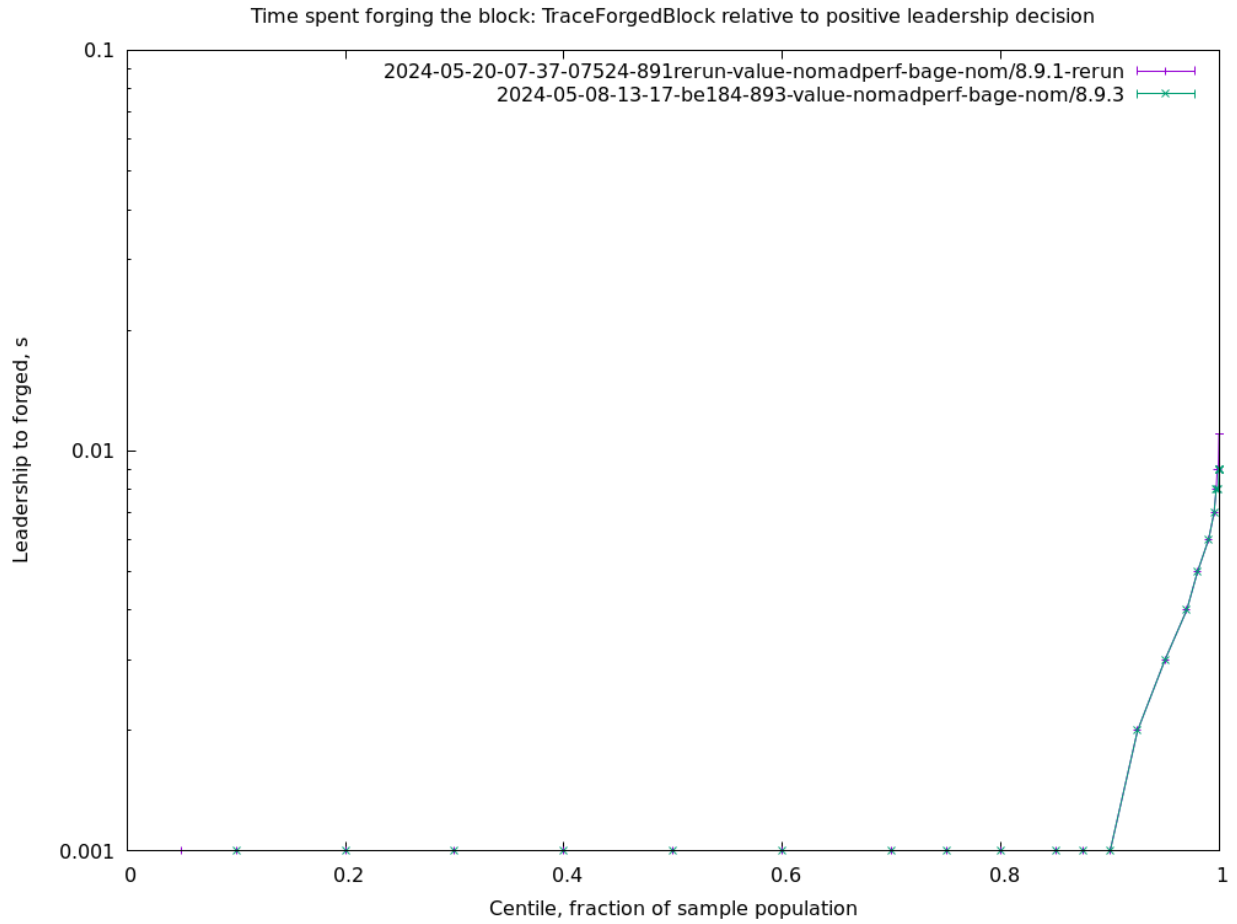
**Ledger ticking (cdfForgerTicked)** Time spent ticking the ledger state (TraceForgeTickedLedgerState), relative to leadership check completion



**Mempool snapshotting (cdfForgerMemSnap)** Time spent taking a mempool snapshot (TraceForgingMempoolSnapshot), relative to ledger ticking conclusion

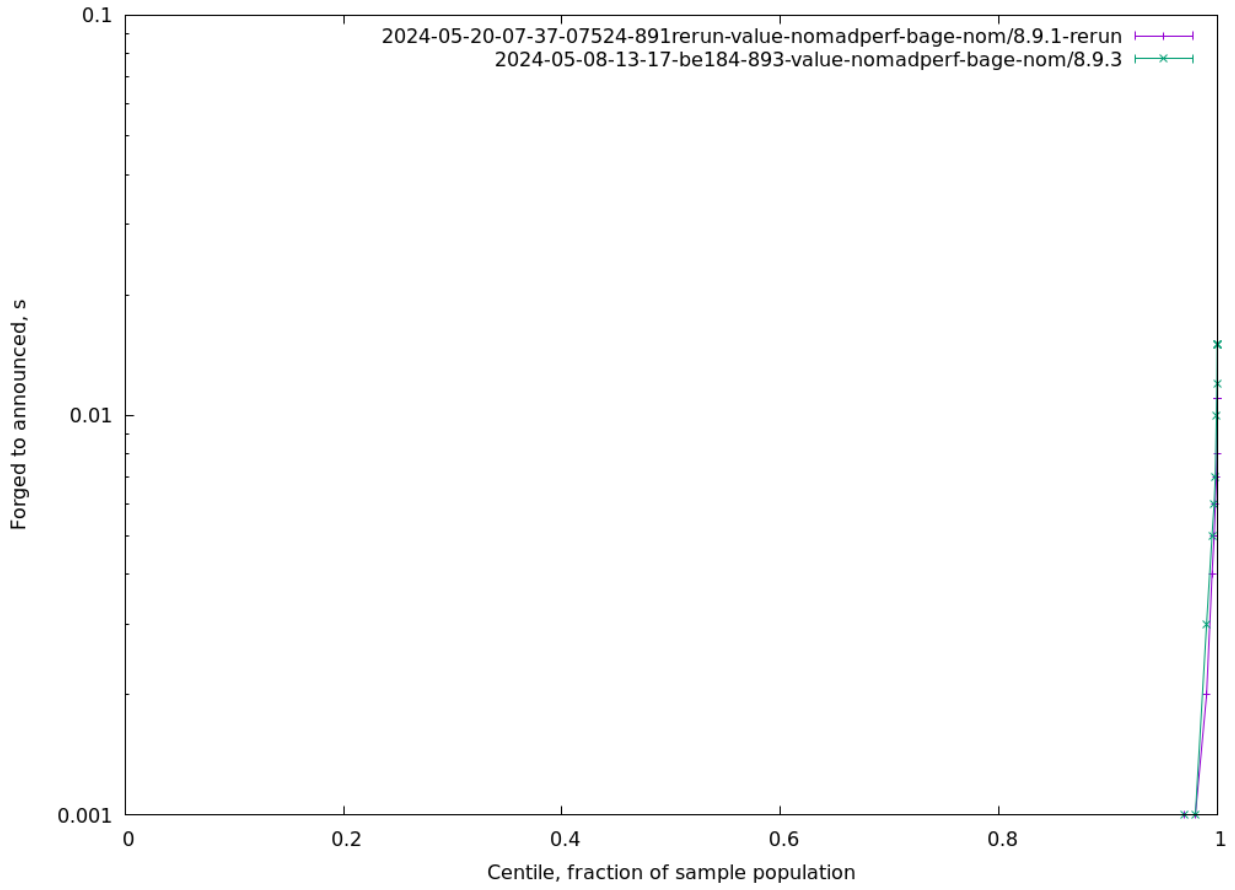


**Leadership to forged (cdfForgerForge)** Time spent forging the block: TraceForgedBlock relative to positive leadership decision

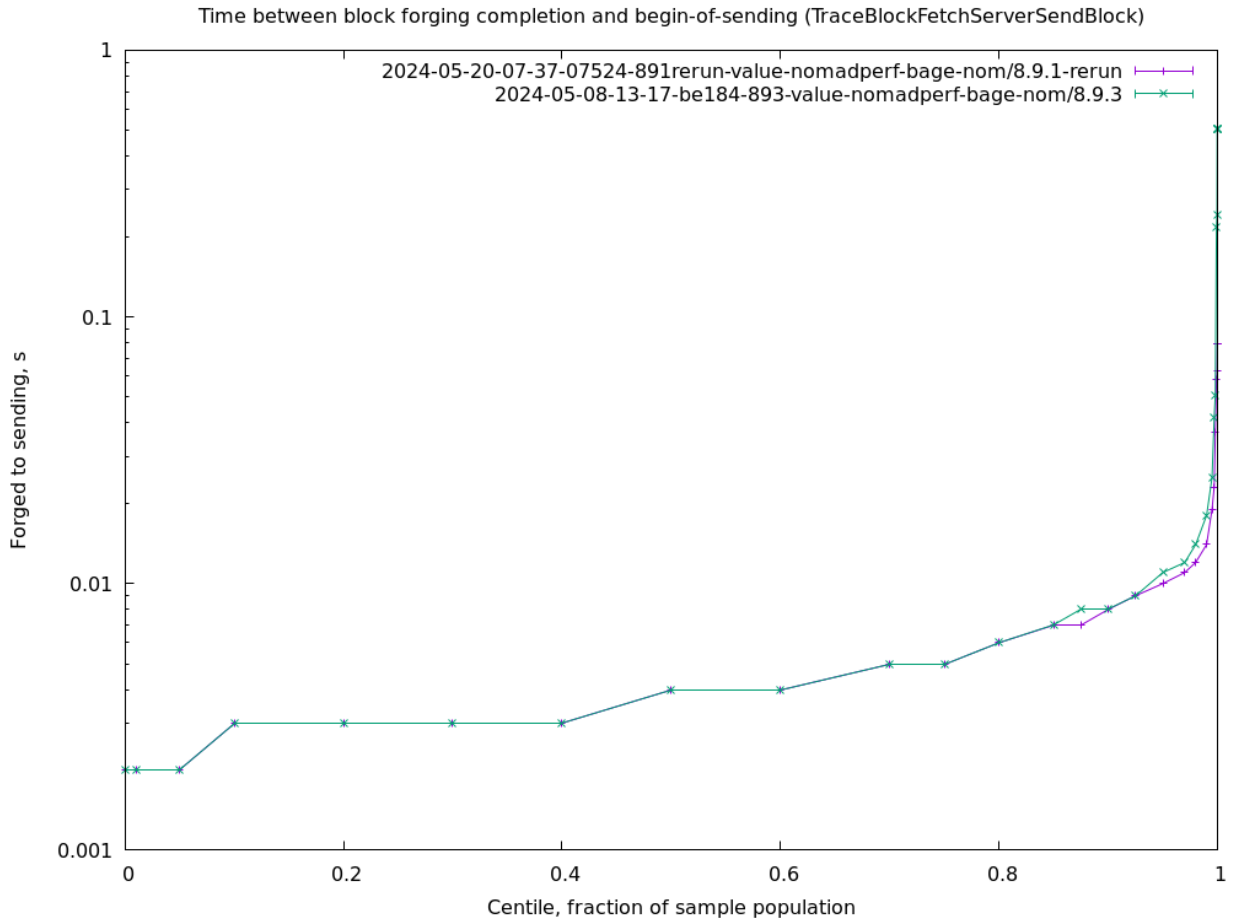


**Forged to announced (cdfForgerAnnounce)** Time between block forging completion and header announcement (ChainSyncServerEvent.TraceChainSyncServerRead.AddBlock)

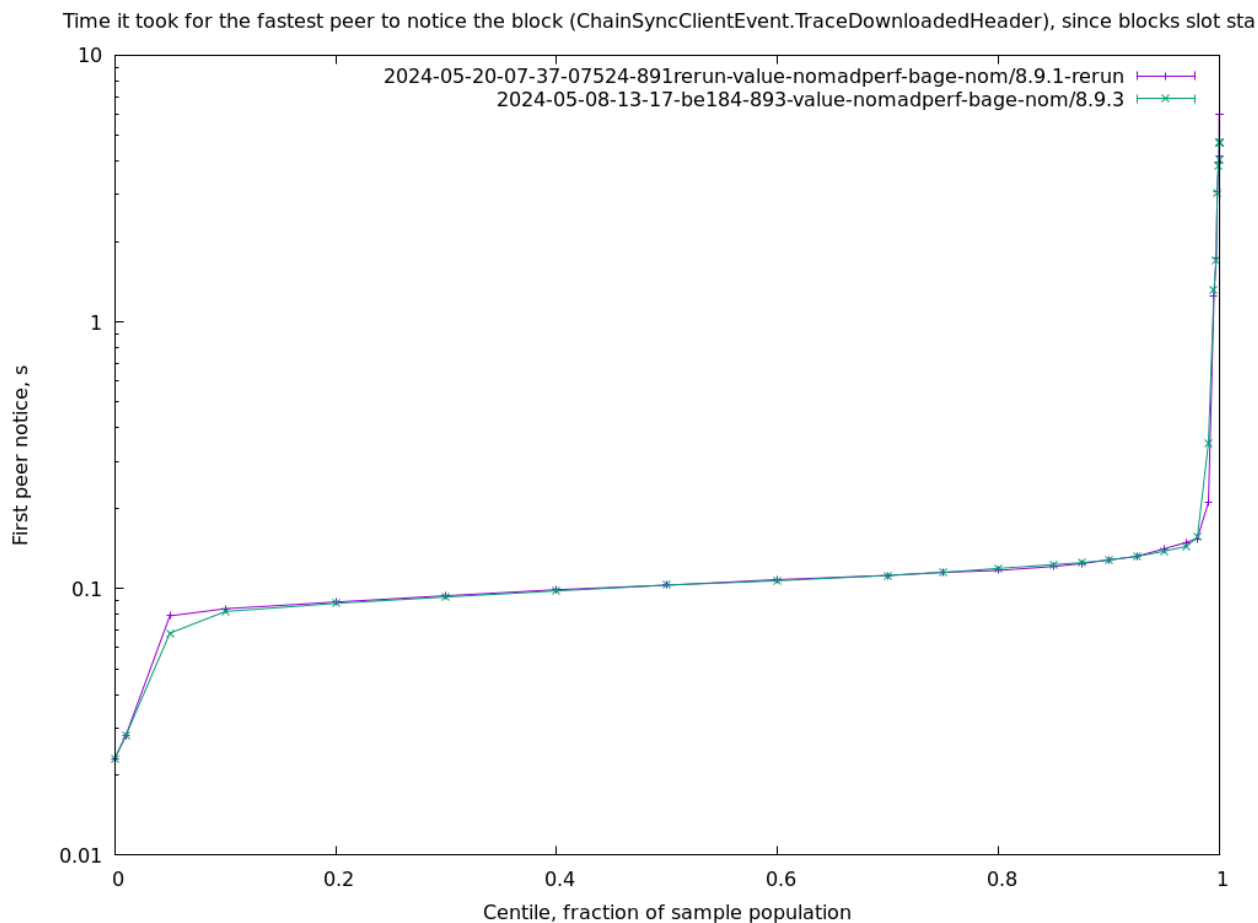
Time between block forging completion and header announcement (ChainSyncServerEvent.TraceChainSyncServerRead.Add)



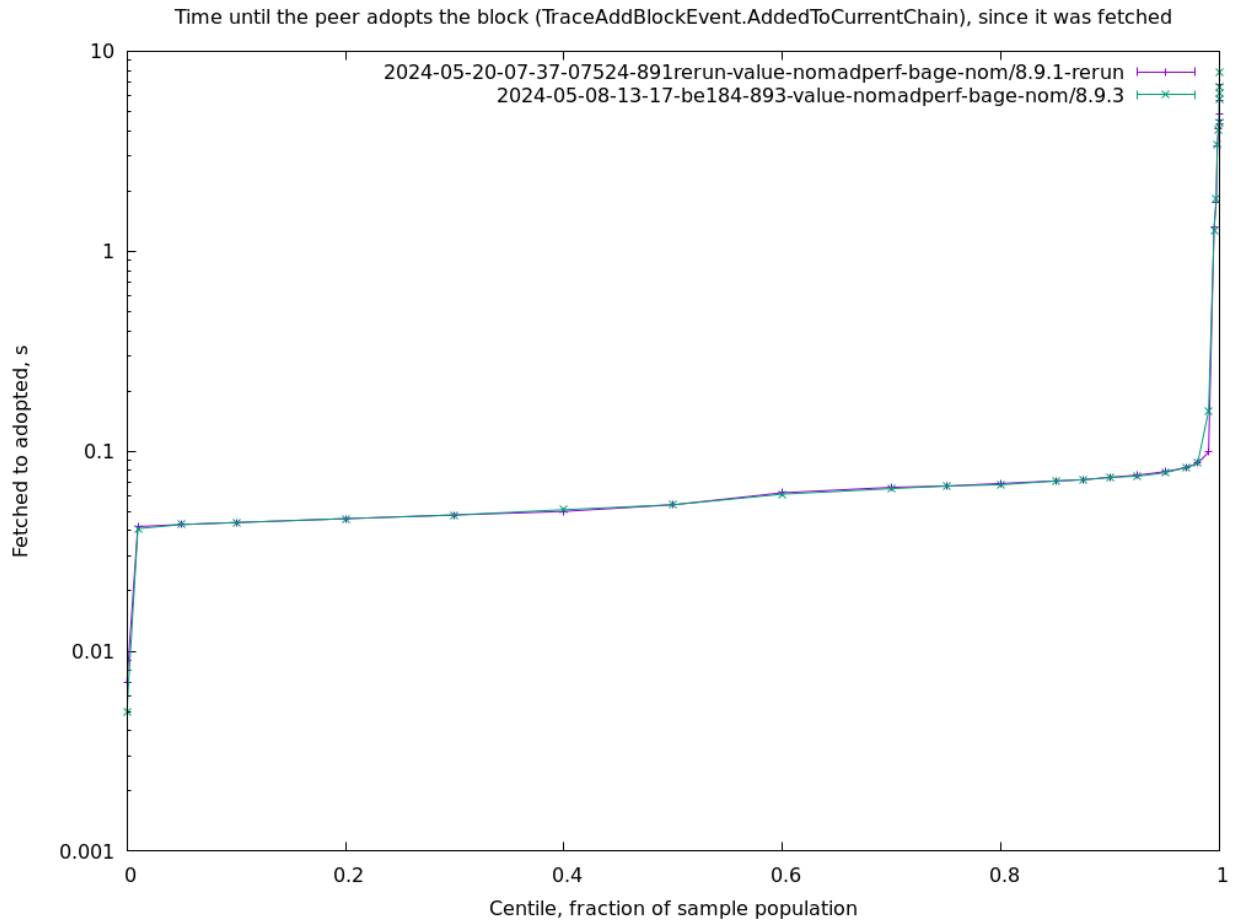
**Forged to sending (cdfForgerSend)** Time between block forging completion and begin-of-sending (TraceBlockFetchServerSendBlock)



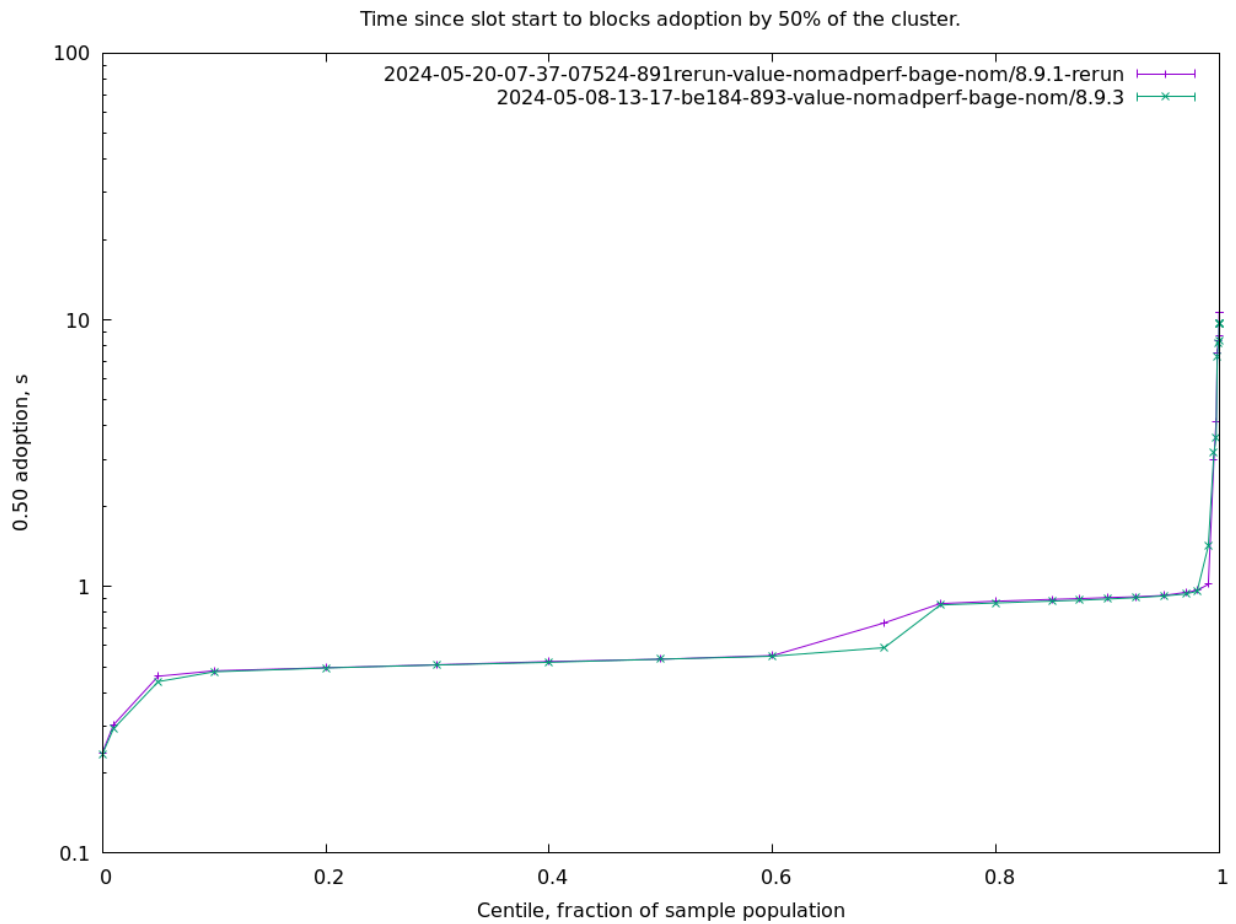
**First peer notice (cdfPeerNoticeFirst)** Time it took for the fastest peer to notice the block (ChainSyncClientEvent.TraceDownloadedHeader), since block's slot start



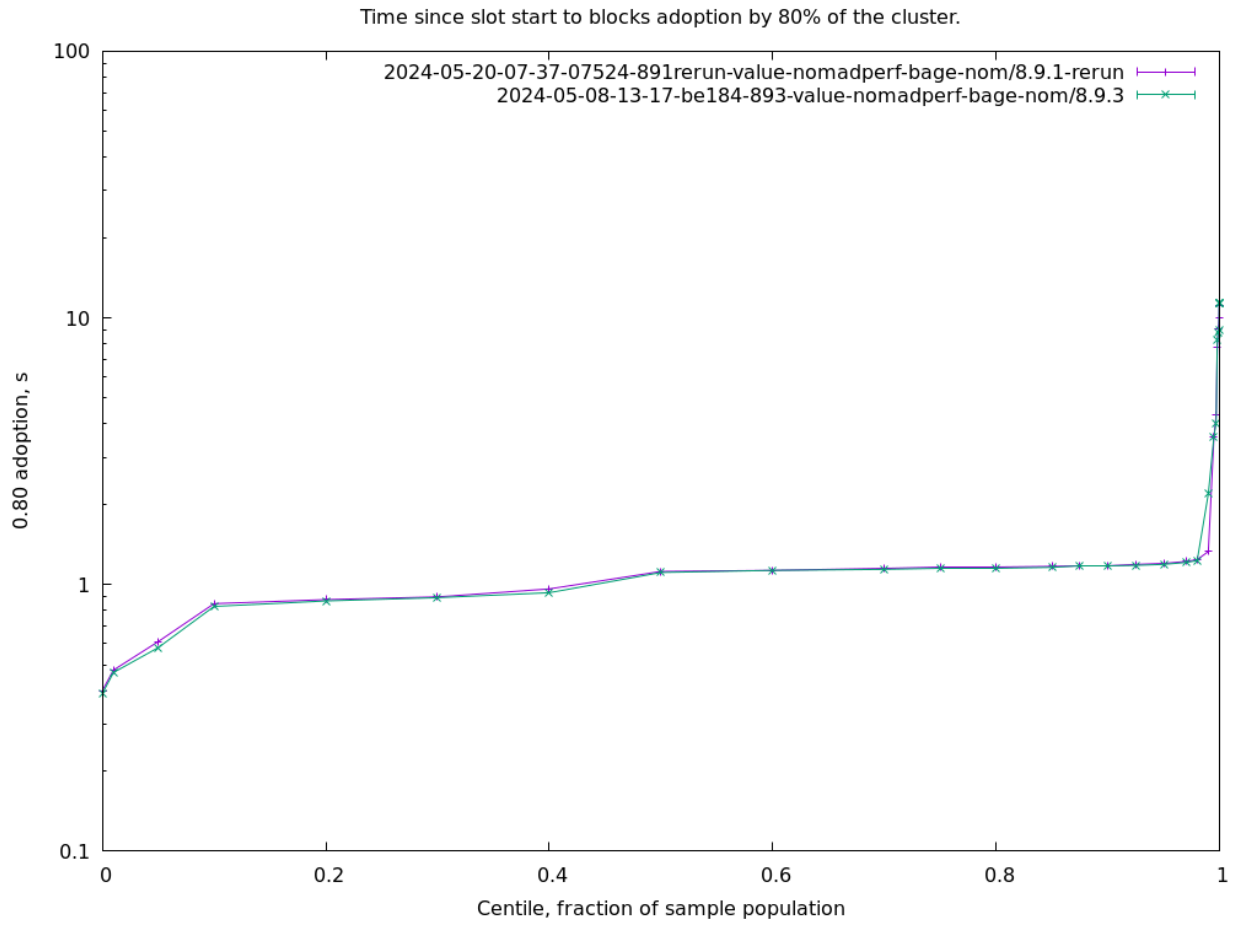
**Fetches to adopted (cdfPeerAdoption)** Time until the peer adopts the block (TraceAddBlockEvent.AddedToCurrentChain), since it was fetched



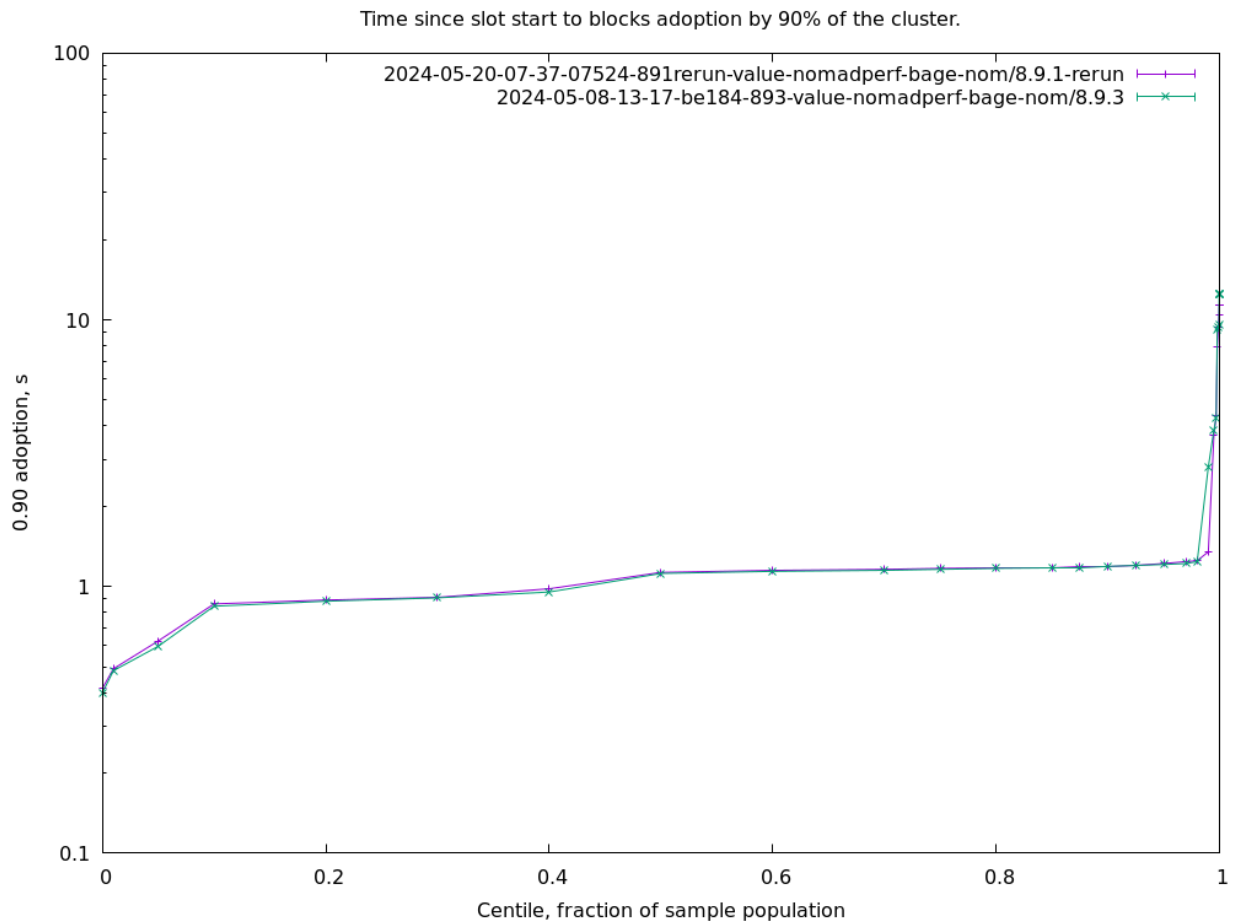
**0.50 adoption (cdf0.50)** Time since slot start to block's adoption by 50% of the cluster.



**0.80 adoption (cdf0.80)** Time since slot start to block's adoption by 80% of the cluster.



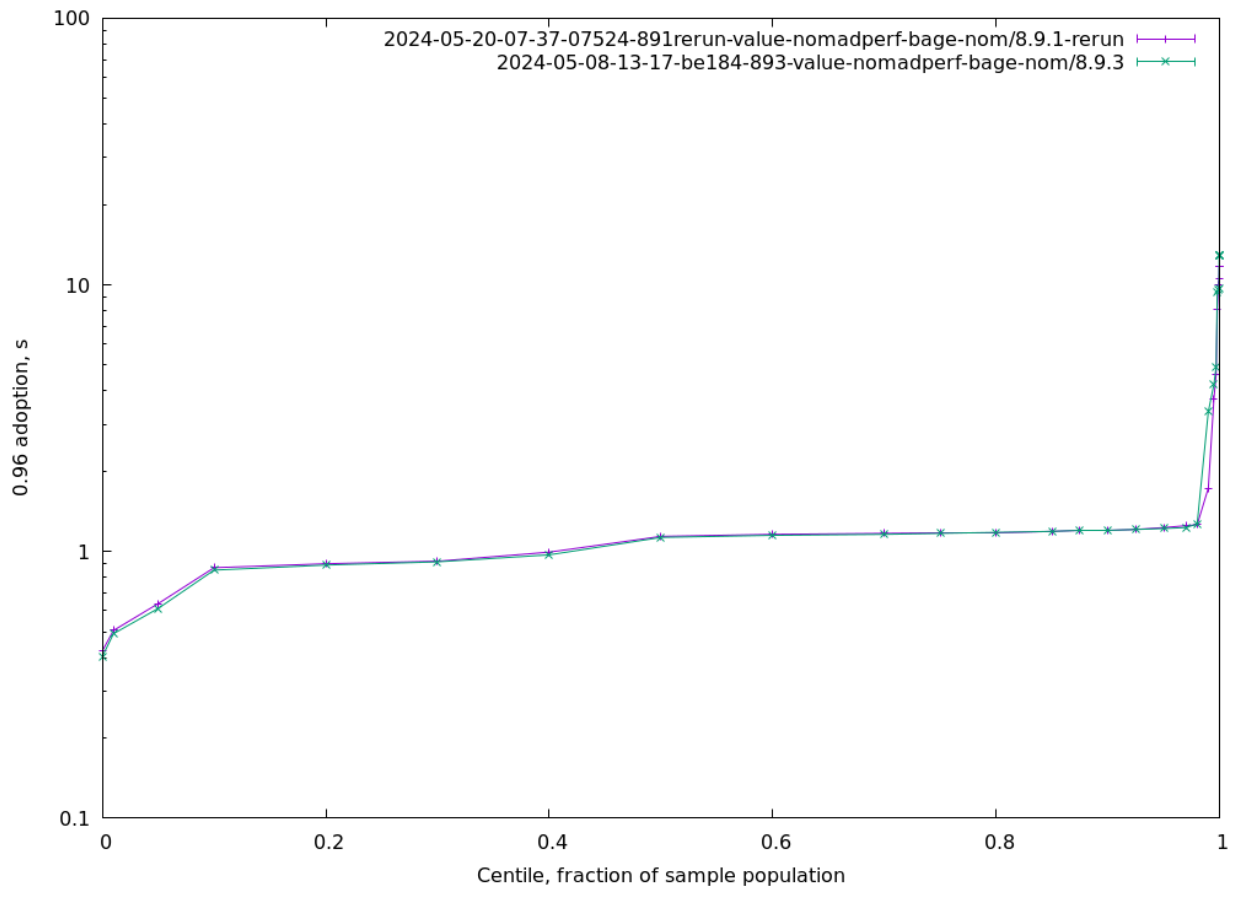
**0.90 adoption (cdf0.90)** Time since slot start to block's adoption by 90% of the cluster.



**0.96 adoption (cdf0.96)** Time since slot start to block's adoption by 96% of the cluster.



Time since slot start to blocks adoption by 96% of the cluster.



## Part II

# Appendix B: data dictionary

# Chapter 4

## Block propagation metrics

**0.50 adoption (cdf0.50)** Time since slot start to block's adoption by 50% of the cluster.

**0.80 adoption (cdf0.80)** Time since slot start to block's adoption by 80% of the cluster.

**0.90 adoption (cdf0.90)** Time since slot start to block's adoption by 90% of the cluster.

**0.92 adoption (cdf0.92)** Time since slot start to block's adoption by 92% of the cluster.

**0.94 adoption (cdf0.94)** Time since slot start to block's adoption by 94% of the cluster.

**0.96 adoption (cdf0.96)** Time since slot start to block's adoption by 96% of the cluster.

**0.98 adoption (cdf0.98)** Time since slot start to block's adoption by 98% of the cluster.

**1.00 adoption (cdf1.00)** Time since slot start to block's adoption by 100% of the cluster.

**Height & slot battles (cdfBlockBattle)** For a given block, number of all abandoned blocks at its block height. Sum of height and slot battles

**Block size (cdfBlockSize)** Block size, in bytes

**Chained to forged block ratio (cdfBlocksChainedRatio)** For each host, ratio of blocks that made into chain / all forged

**Filtered to chained block ratio (cdfBlocksFilteredRatio)** For each host, ratio of blocks that passed filtering / all on chain

**Blocks per host (cdfBlocksPerHost)** For each host, number of blocks made during the entire observation period

**Forged to self-adopted (cdfForgerAdoption)** Time between block forging completion and adoption (TraceAdoptedBlock)

**Forged to announced (cdfForgerAnnounce)** Time between block forging completion and header announcement (ChainSyncServerEvent.TraceChainSyncServerRead.AddBlock)

**Slot start to announced (cdfForgerAnnounceCum)** Time since slot start until header announcement (ChainSyncServerEvent.TraceChainSyncServerRead.AddBlock)

**Acquired block context (cdfForgerBlkCtx)** Block context acquired (TraceBlockContext), relative to forge loop beginning

**Leadership to forged (cdfForgerForge)** Time spent forging the block: TraceForgedBlock relative to positive leadership decision

**Leadership check duration (cdfForgerLead)** Leadership check duration (TraceNodeIsNotLeader, TraceNodeIsLeader), relative to ledger view acquisition

**Acquired ledger state (cdfForgerLgrState)** Ledger state acquired (TraceLedgerState), relative to block context acquisition

**Acquired ledger view (cdfForgerLgrView)** Ledger view acquired (TraceLedgerView), relative to ledger state acquisition

**Mempool snapshotting (cdfForgerMemSnap)** Time spent taking a mempool snapshot (TraceForgingMempoolSnapshot), relative to ledger ticking conclusion

**Forged to sending (cdfForgerSend)** Time between block forging completion and begin-of-sending (TraceBlockFetchServerSendBlock)

**Started forge loop iteration (cdfForgerStart)** Forge loop iteration delay (TraceStartLeadershipCheck), relative to slot start

**Ledger ticking (cdfForgerTicked)** Time spent ticking the ledger state (TraceForgeTickedLedgerState), relative to leadership check completion

**Fetches to adopted (cdfPeerAdoption)** Time until the peer adopts the block (TraceAddBlockEvent.AddedToCurrentChain), since it was fetched

**Fetches to announced (cdfPeerAnnounce)** Time it took a peer to announce the block (ChainSyncServerEvent.TraceChainSync), since it was fetched

**Fetch duration (cdfPeerFetch)** Time it took the peer to complete fetching the block (BlockFetchClient.CompletedBlockFetch), after having requested it

**First peer fetch (cdfPeerFetchFirst)** Time it took for the fastest peer to fetch the block (BlockFetchClient.CompletedBlockFetch), since block's slot start

**First peer notice (cdfPeerNoticeFirst)** Time it took for the fastest peer to notice the block (ChainSyncClientEvent.TraceDownloadedHeader), since block's slot start

**Notice to fetch request (cdfPeerRequest)** Time it took the peer to request the block body (BlockFetchClient.SendFetchRequest), after it have seen its header

**Fetches to sending (cdfPeerSend)** Time until the peer started sending the block (BlockFetchServer.SendBlock), since it was fetched

## Chapter 5

# Cluster performance metrics

**RTS alloc rate (Alloc)** RTS-reported allocation rate, MB/sec

**Process CPU usage (CentiCpu)** Kernel-reported CPU process usage, % of a single core

**RTS GC CPU usage (CentiGC)** RTS-reported GC CPU usage, % of a single core

**RTS Mutator CPU usage (CentiMut)** RTS-reported mutator CPU usage, % of a single core

**Filesystem reads (FsRd)** Number of bytes which this process really did cause to be fetched from the storage layer, per second

**Filesystem writes (FsWr)** Number of bytes which this process caused to be sent to the storage layer, modulo truncate(), per second

**Major GCs (GcsMajor)** Major garbage collection RTS events

**Minor GCs (GcsMinor)** Minor garbage collection RTS events

**RTS heap size (Heap)** RTS-reported heap size, MB

**RTS live GC dataset (Live)** RTS-reported GC live data size, MB

**Network reads (NetRd)** Network reads, kB/sec

**Network writes (NetWr)** Network writes, kB/sec

**Kernel RSS (RSS)** Kernel-reported RSS (Resident Set Size) of the process, MB

**Block context acquisition delay (cdfBlkCtx)** Block context acquired (TraceBlockContext), relative to forge loop beginning

**Interblock gap (cdfBlockGap)** Time between blocks

**Chain density (cdfDensity)** Block/slot ratio, for the last 'k' slots

**Leadership check duration (cdfLeading)** Leadership check duration (TraceNodeIsNotLeader, TraceNodeIsLeader), relative to ledger view acquisition

**Ledger state acquisition delay (cdfLgrState)** Ledger state acquired (TraceLedgerState), relative to block context acquisition

**Ledger view acquisition delay (cdfLgrView)** Ledger view acquired (TraceLedgerView), relative to ledger state acquisition

**CPU 85% spans (cdfSpanLensCpu)** Length of over-85% CPU usage peaks, slots

**CPU spans at Ep boundary (cdfSpanLensCpuEpoch)** Length of over-85% CPU usage peaks, starting at epoch boundary, slots

**Forge loop tardiness (cdfStarted)** Forge loop iteration start delay (TraceStartLeadershipCheck), relative to slot start

**Forge loop starts (cdfStarts)** For any given slot, how many forging loop starts were registered