Black Block Recorder: Immutable Black Box Logging via rosbag2 and DLTs

By: Ruffin White, Gianluca Caiazza





Motivation

With maturing R&D, robots have propagated into many industries and economies, including manufacturing, transport, logistics, and more.

As cyber-physical systems (CPS), the risks from and threats to robots can surpass that incurred by conventional Internet of Things (IoT) devices.

Yet, are the methods used to monitor/audit such CPS devices in pace with robotic proliferation?



Motivation

Event Data Recorders (**EDRs**) in automotive, or Black Box Flight recorders in aerospace industries serve a multitude of critical directives, including:

- Dataset & Model Acquisition
- Performance Benchmarking
- Quality Control and Testing
- Continuous Health Monitoring
- Diagnostic Debugging & Triage
- System Auditing & Verification
- Digital Forensic Investigations





Motivation

With high-stake deployments come high-stake consequences, thus the extraordinary incentives evident to **secure** vs. **circumvent** EDR devices:

- Hold liable parties accountable
- Guide regulatory & legislative policy VS.
- Expunge incriminating evidence
- Conspire to falsify historic events

Directive: to verify the *integrity*, *authenticity*, and *completeness* of robotic logs under the threat of malicious/erroneous <u>insertion</u>, or <u>omission</u>.

The five Robots Death by robot: the new mechanised danger in our changing world

As the use of autonomous machines increases in society, so too has the chance of robot-related fatalities



Related Work

- [1] Jordi Cucurull, and Jordi Puiggalí. "Distributed immutabilization of secure logs." International Workshop on Security and Trust Management. Springer, Cham, 2016.
- [2] Andrew Sutton, and Reza Samavi. "Blockchain enabled privacy audit logs." International Semantic Web Conference. Springer, Cham, 2017.

1] Table 1. Data immutabilized within the transaction output

Name	Bytes	Content	Description
prefix	2	"SL"	OP_RETURN prefix
version	1	1	Version of the data structure
logId	16	-	Unique log identifier
H_{j}	32	_	Hash of the checkpoint



[3]

Related Work

- [3] Sebastian Taurer, Bernhard Dieber, and Peter Schartner.
 "Secure data recording and bio-inspired functional integrity for intelligent robots." 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2018.
- [4] Viktoras Kabir Veitas, and Simon Delaere. "In-vehicle data recording, storage and access management in autonomous vehicles." arXiv preprint arXiv:1806.03243 (2018).





Constraints

CPSs such as commercially viable mobile robots incur contains in addition to just cyber security:

- Limited energy capacity
 - Efficiency is key to prolong deployment
 - Minimize computational overhead
- Build of Materials
 - Practical costs and availability at scale
 - Consumer grade hardware compatible
- Finite memory storage
 - Event data recording must be continuous
 - Maximize data retention for auditability
- Unreliable network connectivity
 - Uplink bandwidth a precious commodity
 - Flexible QoS to minimize overhead











Roles & Obligated Parties

Custodian: Robot or autonomous vehicle OEM Reporter: Trusted Logger or Recorder Enclave Owner: End-User or Operator Auditors: Regulatory Agencies or Gov. Attacker: Malicious Actor





Party: Custodian

Obligated subject of log content and tasked with log preservation. e.g. Robot or autonomous vehicle OEM.



Party: Reporter

An independent party responsible for faithfully recording events. e.g. Trusted Logger or Recorder Enclave.



Party: Owner

Mediating party that has a stake in ensuring log integrity, authenticity, confidentiality. e.g. End-User or Operator.



Party: Auditor

Observing parties called upon to investigate and validate record archives. e.g. Regulatory Agencies or Gov.



Party: Attacker

Any party intent on modifying the record to falsify events. e.g. Malicious actor expunging incriminating evidence.



Storage Requirements

1. Data provision conditions

requires consent on behalf of the **Owner** who transitively controls the log assets tracked.

- 2. Fair and undistorted competition trust should be distributed and shared across all validators (a.k.a **Custodians**).
- 3. Data privacy and data protection the co-location of logs external to that of the **Custodian** must be prevented.
- 4. Tamper-proof access and liability integrity and authenticity of logs must derive from an independent **Reporter**.
- 5. Data availability economy

health and transparency of logs are contingent upon giving **Auditors** appropriate access.











System Properties

- Secure identification of physical data sources attestation between devices trusted by the Custodian and Reporter.
- 2. Metadata enrichment

log event context may be associated to respective **Owner**, **Custodian** and **Reporter** parties.

- 3. Data exchange and messaging authenticated encryption is used in establishing secure connectivity between parties.
- 4. Data recording & storage

reporting remains flexible in terms of QoS as well as reasonable in resource consumption.

5. Access management

rights, obligations, and authorization of parties must be explicitly defined and enforceable.

FOREIGN









Approach

- Introducing Black Block Recorder (BBR)
- To mediate separation of concerns, only the recording process need be enclaved
- Enclaved process generates the logs by capturing broadcasted message traffic
- Before streaming logs to arbitrary storage, data must first be rendered immutable



Approach

- Striding checkpoints are disseminated to the a blockchain for eventual log auditing
- Checkpoints are comprised of cryptographic linked integrity proofs that are indexed
- Thus we explore EDRs based upon Distributed Ledger Technology (DLT)



 $Chk_{i} = (i, h_{i}) \quad h_{i} = HMAC(h_{i-1}, LogMsg_{i})$ where $h_{0} \leftarrow \{0, 1\}^{m}$ (1)

Information Flow

- While a robot platform is held suspect, an enclave (e.g. TTE) is reserved for the reporter
- Event data is received within the enclave via SROS2 using authenticated encryption
- Linked integrity proofs for each asset tracked are batched and signed via BBR rosbag plugin
- Batches are then signed within the enclave, then co-signed by the robot's via BBR bridge
- Thus an append-only forgery to DLT record would necessitate the collusion of both the custodian and its respective assigned reporter





Event Data → Enclave → ROS2 Bag Plugin → Bridge → DLT Validator → DLT Network

Schema Structure

- Recorded event data is structured using ROS2bag files and compatible SQLite schema
- Data insertions into the database are achieved via 2D-array hash-chains of checkpoints
- Where the primary axis checkpoints each bagfile topic's genesis-block and meta-info
- While the secondary axis checkpoints the insertion of respective message data
- This coupling affords a holistic integrity proof of the entire database, while preserving topics as a time series atomic, fostering parallelism





Topics Schema

- The nonce and digest for each topic are retained to verify even IDL info & metadata
- Here we'll choose to hash the topic name metadata separate from the IDL info
- This allows for message data to remain verifiable even when want to remap topics
- As a superset of the default SQLite schema, bags from the BBR plugin remain compliant



Topics					BBR		
	id	name	type	format	nonce	digest	
	1	/foo	gps	rtps	bits _A	bits _{A0}	
	2	/bar	imu	rtps	bits _B	bits _{B0}	
	3	/baz	lidar	rtps	bits _c	bits _{co}	

$bits_{bag} \leftarrow \{0,1\}^m$	(2)
$bits_A \leftarrow HMAC(bits_{bag}, Proto(name_{bag}))$	(3)
$bits_{A_0} \leftarrow HMAC(bits_A, Proto(type_A, format_A))$	(4)
$bits_B \leftarrow HMAC(bits_{A_0}, Proto(name_A))$	(5)
$bits_{B_0} \leftarrow HMAC(bits_B, Proto(type_B, format_B))$	(6)
$bits_C \leftarrow HMAC(bits_{B_0}, Proto(name_B))$	(7)
$bits_{C_0} \leftarrow HMAC(bits_C, Proto(type_C, format_C))$	(8)

Messages Schema

- Each message is checkpointed by the time of arrival and data payload received
- Here we'll choose not to hash the topic ID, given message is bound to the topic digest

tabase Str	icture Browse	Data Edit Pra	gmas Execut	e SQL	Edit Datab	ase cell				
le: 🔳 me	ssages 🔹 😂 🕥	6	New Record	Delete Record	Mode: Bi	nary -			Import	Export Set as N
id	topic_id	timestamp	data	bbr_digest	0000 40) 55 d9	7a ac 40 0	d 16 22 62	31 b2 86 cf 29 c3 20 ec 60 81 49 c2	@UUz –@ " b1² []ĭ 8ávî Å: □□(C) ``
Filter	Filter	Filter	Filter	Filter	0020					say in a blirty i
1	1	156295405	BLOB	@U\$2\$@8						
2	1	156295405	BLOB		Type of da	ta curre	ntly in cell: Bi	narv		
3	1	156295405	BLOB	AAI'AAAI'''	32 byte(s)					Ap
4	1	156295405		VV/VØV						
5	1	156295405		10 442145	DB Schema					
6	1	156295405	BLOB	VVIVV ~I	Name			Туре	Schema	
7	1	156205405	RLOP	5CWJWWYQ	🝷 🗏 Tabl	es (2)				
/	1	150295405		V.> LVV4	- n 🗆 -	nessages			CREATE TABLE n	nessages(id INTEGE.
8	1	150295405		VPWHOWW		a Id		INTEGER	Id INTEGER	
9	1	156295405				timest	0	INTEGER	`timestame` INT	ECER NOT NULL
10	1	156295406	BLOB	at array Born		data	amp	BLOB	'data' BLOB NO	
						bbr di	aest	BLOB	`bbr_digest` BLC	B NOT NULL
					* 🗉 b	opics	J		CREATE TABLE t	opics(id INTEGER P.
						id 😡		INTEGER	'id' INTEGER	
					6	name		TEXT	'name' TEXT NC	T NULL
					6	type		TEXT	'type' TEXT NOT	NULL
						serializ	ation_forma	TEXT	`serialization_fo	rmat' TEXT NOT N
					6	bbr_nc	once	BLOB	'bbr_nonce' BLC	B NOT NULL
						bbr di	nest	BLOB	'bbr digest' BLC	B NOT NULL

Messa	BBR			
id	topic	time	data	digest
1	id 1	int64	bits	bits _{A1}
2	id 2	int64	bits	bits _{B1}
3	id 1	int64	bits	bits _{A2}

 $bits_{A_1} \leftarrow HMAC(bits_{A_0}, Proto(time_{A_1}, data_{A_1}))$ (9) $bits_{B_1} \leftarrow HMAC(bits_{B_0}, Proto(time_{B_1}, data_{B_1}))$ (10) $bits_{A_2} \leftarrow HMAC(bits_{A_1}, Proto(time_{A_2}, data_{A_2}))$ (11)

- This allows for message data to remain verifiable even when want to splice bags
- Checkpoints link to previous message of the same topic, not necessarily previous insert

- Checkpoints ensure immutability, but not authenticity and non-repudiation by itself
- Smart Contracts (SC) encapsulate access control logic for state in permissioned DLTs
- Digital Asset Modeling Language (DAML) is used to formalize contract specification
- Record/Checkpoint data structures define the base abstractions used in SC modeling



16 data Record = Record with -- data type struct -- "/sensors/exteroceptive/qps" 17 r name: Text r type: Text -- "qps" 18 19 r format: Text -- "rtps" 20 r nonce: Text -- "<bits A>" 21 r digest: Text -- "<bits A 0>" r checkpoints: [Checkpoint] -- monotonic list 22

42 data Checkpoint = Checkpoint with -- data type struct 43 c_proof: Text -- "<bits_A_1>" 44 c_stamp: Int -- 1



- An EDR SC is formed from parties involved voluntarily entering as obligated signatories
- External parties may also provisioned observational access to the SC state
- Control for creating new records, to retain checkpoints, is delegated to the reporter

```
template Edr -- Smart Contract for EDRs
2
   with
3
      auditors: [Party] -- Regulatory Agencies
      custodian: Party
                        -- Robot/Vendor Identity
4
                        -- User/Operator Identity
5
      owner: Party
6
      reporter: Party
                        -- Logger/TEE Identity
7
    where
8
      signatory custodian, owner, reporter -- obligated
9
      observer auditors -- non-obligated parties
10
      ensure unique (custodian :: owner :: reporter)
11
      controller reporter can -- create many Records
12
        non consuming Edr Record : ContractEd Edr Record
          with record: Record
13
14
          do create Record with edr = this; record
```



- An Record SC is formed from a reference to the EDR SC and initial record struct object
- Only the **Reporter** is permitted the choice to append checkpoints into the ledure state
- Valid checkpoints are appended by asserting stamps remain monotonically increasing
- Both **Reporter** and **Owner** are permitted the choice to permanently finalize the record



```
23 template Edr Record -- Smart Contract for Records in EDR
24
   with
                     -- Reference EDR of origin
25
     edr: Edr
26
     record: Record -- Initial Record state
27
   where
28
      signatory edr.owner, edr.reporter
29
      observer edr. auditors -- custodian can be excluded
      choice EdrRecord Append : ContractEd Edr Record
30
31
          with checkpoints: [Checkpoint] -- [] for batching
          controller edr.reporter -- Only Reporter appends
32
33
          do let -- Update Record with added checkpoints
              is valid = check Monotonic record checkpoints
34
35
              record = append Checkpoints record checkpoints
36
            assert (is valid == True) -- Error on invalid
37
            create Record with edr; record = record
38
      choice EdrRecord Finalize : () -- Archives Contract
39
        controller edr.owner, edr.reporter
          do return () -- Finalized Record is un-appendable
40
```



- An EDR SC is preliminarily proposed via pending contract used to collect the necessary multi-party signatories.
- The pending SC is consumed to create the agreed EDR SC
- The **Reporter** may create multiple referencing records and append checkpoints, that itself or the **Owner** may finalize





Results

28

- Drop Rate & Load performance comparison of BBR storage/bridge plugins with regard to ROSBag2's default SQLite plugin.
- Benchmarked via ROS2 Crystal, 2.6GHz Intel i7-6700HQ, with RTI Connext RMW on loopback interface.

Payload	Size (bytes)	Requirements
Chk _i Signed Transaction Signed Batch	$36 \ge 629 \ge 965$	Ledger Disk Storage Network Bandwidth Network Bandwidth

 TABLE I: BBR Payload Allocations

Msg. Size vs Drop Rate @ fixed 1kHz Msg. Freq.



Conclusion

- Engineering trade offs always exist
 - Security vs. Performance
 - Immutability vs. Drop-Rate
- However trade offs can be balanced
 - Adjusting QoS for sensitive topics
 - Sparsity via Striding checkpoint
- Promising uses of DLTs in Cryptobotics
 - EDR applications leveraging distributed consensus
 - When Byzantine Fault Tolerance is advantageous

Mes	sages	BBR		
id	topic		digest	
1	id 1		bits _{A1}	->
2	id 1		bits _{A2}	
3	id 1		bits _{A3}	
2	id 1		bits _{A4}	->



Future Work

- Benchmark with consensus
 - Profile validator overhead
 - Characterize scale limitations
- Explore alternative DLT pardimes, e.g:
 - Directed Acyclic Graphs vs Blockchains
 - E.g. IOTA's Tangle
- Record at the middleware level (DDS)
 - Capture all traffic like services/actions
 - E.g. Parameter updates, planner feedback

DDS





Acknowledgements





Advances in Autonomous, Distributed and Pervasive systems



Links

32

Current Slides: <u>roscon.ros.org/2019/presentations/</u> <u>ROSCon2019_Black_Block_Recorder.pdf</u>



O GitHub

dledr/bbr_ros2

Black Block Recorder with ROS2 | Immutable Logging via Blockchain - dledr/bbr_ros2

Ç GitHub

• ¬ ros2/sros2

tools to generate and distribute keys for SROS 2. Contribute to ros2/sros2 development by creating an account on GitHub.

Hyperledger



Hyperledger Sawtooth – Hyperledger

A modular platform for building, deploying, and running distributed ledgers.

Hyperledger



Hyperledger – Open Source Blockchain Technologies

Hyperledger is a multi-project open source collaborative effort hosted by The Linux Foundation, created to advance cross-industry blockchain technologies.

Ç GitHub

digital-asset/daml



The DAML smart contract language. Contribute to digital-asset/daml development by creating an account on GitHub.

Questions?

- A113: Full Autopilot Override
- BBR: Black Block Recorder
- BFT: Byzantine Fault Tolerance
- CPS: Cyber-Physical System
- DAG: Directed Acyclic Graph
- **DFI**: **D**igital Forensic Investigation
- **DLT**: Distributed Ledger Technology
- DSA: Digital Signature Algorithm
- EDR: Event Data Recorder
- MAC: Message Authentication Code ROS: Robot Operating System
- 33 SROS: Secure ROS



References



R. White, G. Caiazza, A. Cortesi, Y. I. Cho and H. I. Christensen, "Black Block Recorder: Immutable Black Box Logging for Robots via Blockchain," in *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3812-3819, Oct. 2019. doi: <u>10.1109/LRA.2019.2928780</u>

"WALL-E", Pixar Animation Studios, Walt Disney Studios Motion Pictures, (film) Jun. 2008.

M. Luthi, A. Sparrow, et al. "WALL-E", in BOOM! Studios, (comic) Nov. 2009 - Jun. 2010.

More about:

34

Ruffin: <u>about.me/ruffin</u> Gianluca: about.me/caiazza

