PLANCKCHRON FLAGSHIP TECHNICAL SERIES

RESILIENT SYSTEMS / VERSION 1.0 / 09 AUGUST 2026

EchoNet

Resilient Coordination Across Intermittent and Partitioned Networks

ABSTRACT

EchoNet is presented as a prototype-stage coordination architecture for environments where end-to-end connectivity cannot be assumed. The design combines local-first state, priority-aware store-carry-forward messaging, delay-tolerant interoperability, explicit conflict handling, device identity, bundle-level security, and recovery reconciliation. It does not claim global coverage, a deployed network, or validated latency and throughput. The paper defines message and node contracts, four operating modes, a threat model, an evaluation matrix, and a phased prototype plan. It draws from IETF Bundle Protocol standards, delay-tolerant networking architecture, public space-communications experience, and IoT device security guidance. Those sources establish prior art and design constraints; they do not establish Planckchron performance or institutional affiliation.

PLANCKCHRON RESEARCH

Self-published technical paper / not peer-reviewed / open access

planckchron.com/publications/echonet-resilient-intermittent-network-coordination

TECHNICAL NOTE

Executive brief

EchoNet is presented as a prototype-stage coordination architecture for environments where end-to-end connectivity cannot be assumed. The design combines local-first state, priority-aware store-carry-forward messaging, delay-tolerant interoperability, explicit conflict handling, device identity, bundle-level security, and recovery reconciliation. It does not claim global coverage, a deployed network, or validated latency and throughput. The paper defines message and node contracts, four operating modes, a threat model, an evaluation matrix, and a phased prototype plan. It draws from IETF Bundle Protocol standards, delay-tolerant networking architecture, public space-communications experience, and IoT device security guidance. Those sources establish prior art and design constraints; they do not establish Planckchron performance or institutional affiliation.

EVIDENCE BOUNDARY

Prototype design only. No deployed network, global coverage, guaranteed delivery, validated latency, throughput, range, security certification, or operational mission capability is claimed.

Four operating conclusions

- Intermittency is an operating condition, not an exception to hide behind retries.
- The message needs its own identity, priority, expiry, security, provenance, and conflict policy.
- Local autonomy and later reconciliation must be designed together.
- Evaluation must expose the delivery-latency-energy-security trade space under controlled disruptions.

STATUS	**Prototype design**
PUBLISHER	Planckchron Inc.
CITATION	*Planckchron Research (2026). EchoNet: Resilient Coordination Across Intermittent and Partitioned Networks. Version 1.0.*

TECHNICAL NOTE

Reference architecture

The architecture below is a proposed control and evidence flow. It is a design artifact, not a representation of deployed capability.

01

CREATE

Form a signed message envelope with purpose, priority, expiry, provenance, policy, and destination semantics.

02

QUEUE

Store locally with quota, deduplication, custody state, and resource-aware scheduling.

03

CARRY

Retain useful data through partitions instead of assuming a simultaneous path.

04

FORWARD

Select contacts and routes using link opportunity, priority, trust, energy, and congestion state.

05

DELIVER

Verify integrity, authorization, freshness, and application acceptance.

06

RECONCILE

Resolve duplicate, stale, and conflicting state as connectivity returns.

TECHNICAL NOTE

1. Problem definition

Most application protocols are optimized for a connected path. In disaster response, remote operations, mobile peer groups, damaged infrastructure, and long-delay environments, the path may not exist when a message is created. Links can be asymmetric, scheduled, low-bandwidth, expensive, or short-lived. A resilient system must preserve intent and evidence while waiting for future contact.

EchoNet is proposed as an application and coordination layer that can operate locally, exchange information opportunistically, interoperate with delay-tolerant protocols, and reconcile state after a partition. It is not proposed as a replacement for the public internet or as a claim of a current communications service.

TECHNICAL NOTE

2. Design principles

The architecture uses six principles. Local utility should continue without a central service. Messages should be useful and verifiable when separated from their origin. Expiry and priority should be explicit. Security should operate at the message or bundle layer in addition to link protection. Resource use should be bounded on constrained nodes. Reconnection should trigger deliberate reconciliation rather than silent last-writer-wins behavior.

These principles reflect the delay-tolerant networking view that eventual end-to-end delivery can be built across heterogeneous constituent networks using store-and-forward overlays. EchoNet adds an application-level state and policy model around that foundation.

TECHNICAL NOTE

3. Reference architecture

Each node contains an identity and key store, local event store, policy engine, priority scheduler, contact and route manager, protocol adapters, conflict resolver, and audit telemetry. Applications write typed events rather than raw transport messages. The scheduler converts eligible events into envelopes and selects forwarding opportunities based on destination, priority, expiry, trust, storage, energy, and congestion.

A Bundle Protocol adapter provides a standards-based path for delay-tolerant exchange where supported. Other adapters may use local wireless or conventional IP transports. The coordination layer preserves one semantic contract across transports. A recovery controller performs inventory exchange, deduplication, missing-event requests, conflict surfacing, and convergence reporting after partitions heal.

TECHNICAL NOTE

4. Message contract

A resilient message needs more context than a destination and payload. The proposed envelope includes message identifier, origin identity, creation time, expiry, priority, content type, content hash, confidentiality and integrity policy, destination semantics, dependency identifiers, custody preference, hop and replication limits, conflict class, acknowledgement policy, and provenance chain.

The conflict class is important. Some events are append-only observations. Some are replaceable state. Some are commands whose freshness and authority must be verified at execution time. Some require human review when competing versions exist. Treating all of them as interchangeable records creates silent and potentially unsafe behavior.

TECHNICAL NOTE

5. Operating modes

Connected mode uses preferred links and can synchronize rapidly, while still writing the durable local record. Degraded mode reduces payloads, prioritizes critical events, and increases tolerance for delayed acknowledgement. Partitioned mode preserves local workflows, applies local policy, and queues messages without pretending that global state is current. Recovery mode exchanges summaries, identifies gaps, verifies freshness and authority, and reconciles conflicts before normal operation resumes.

Mode transitions should be observable and testable. Applications need to know when information may be stale, when an action is only locally accepted, and when a remote acknowledgement is still pending.

TECHNICAL NOTE

6. Threat model and safeguards

The initial threat model includes device impersonation, message tampering, replay, malicious or accidental flooding, metadata leakage, compromised nodes, stale command execution, route manipulation, unauthorized forwarding, key loss, and physical capture. Partitioned operation also creates a governance risk: a node may act on incomplete or obsolete policy.

Proposed controls include unique device identity, signed envelopes, encrypted payloads where required, freshness and replay checks, least-privilege application permissions, quotas, rate and replication limits, tamper-evident local logs, key rotation and revocation propagation, secure update capability, protected reset, and explicit safe behavior when policy cannot be refreshed. Bundle Protocol Security is a relevant standards reference, but implementation still requires a system-specific security design and testing program.

TECHNICAL NOTE

7. Evaluation matrix

Evaluation should use a controlled network emulator and real-device pilot nodes. The matrix varies node count, mobility, contact duration, bandwidth, packet loss, asymmetry, partition length, storage pressure, energy budget, clock error, malicious traffic, and node compromise. Workloads should include observations, bulk data, urgent alerts, commands, and conflicting updates.

Primary measures are delivery probability before expiry, end-to-end delay distribution, overhead per delivered useful byte, energy per delivered message, queue stability, duplicate rate, recovery convergence time, conflict rate, stale-action prevention, security overhead, and containment after compromise. Results must state workload and conditions; there is no single meaningful EchoNet performance number.

TECHNICAL NOTE

8. Prototype path and limitations

Phase one implements signed local events, queueing, expiry, and deterministic reconciliation on a small emulator. Phase two adds intermittent links, congestion, and failure injection. Phase three integrates a standards-based delay-tolerant adapter and security blocks. Phase four introduces selected physical devices and an independent protocol and threat review.

Open questions include usable key recovery in disconnected environments, privacy-preserving routing metadata, safe command semantics, fair scheduling across organizations, policy freshness, interoperability testing, and predictable energy behavior. The prototype should not be represented as a public safety, defense, emergency, or space communications capability without specific evidence and authorization.

TECHNICAL NOTE

References and source notes

References are provided for the external standards, public guidance, and primary research that informed this paper. A citation does not imply endorsement, partnership, certification, or validation of Planckchron capability.

01 **RFC 9171**

Bundle Protocol Version 7

https://www.rfc-editor.org/info/rfc9171/
External IETF proposed standard and prior art.

02 **RFC 9172**

Bundle Protocol Security

https://www.rfc-editor.org/info/rfc9172/
External IETF proposed standard and security reference.

03 **RFC 4838**

Delay-Tolerant Networking Architecture

https://www.rfc-editor.org/info/rfc4838/
External research architecture and prior art.

04 **External research reference - NASA DTN Overview**

Delay/Disruption Tolerant Networking

https://www.nasa.gov/communicating-with-missions/delay-disruption-tolerant-networking/
External public research and deployment context; no Planckchron affiliation or capability claim is implied.

05 **NISTIR 8259A**

IoT Device Cybersecurity Capability Core Baseline

https://csrc.nist.gov/pubs/ir/8259/a/final
External public device-security guidance.

PUBLICATION DISCLOSURE

These self-published technical papers are not peer-reviewed, audited, or independent validation. They describe research methods, proposed architectures, and evaluation plans. Product capability exists only where a dated evidence record expressly reports a result. Prototype design only. No deployed network, global coverage, guaranteed delivery, validated latency, throughput, range, security certification, or operational mission capability is claimed.