DTN for Environmental Monitoring

Scott Burleigh, IPNSIG

Environmental Monitoring

- Environmental change affects our lives, and not all of those effects are beneficial.
- But we can take action to mitigate those effects if we can detect, analyze, understand, and even anticipate environmental change.
- Over time, broad deployment of digital instrumentation enables us to detect even very gradual changes in our environment. But the information produced by those instruments is only useful if it can be communicated to research centers for analysis.
- As the scope of this instrumentation increases, the volume of required communication grows. Sustaining this growth is only possible if the management of that communication is automated. Network support is needed.

Internet-supported Instrumentation

- Much of our instrumentation of the environment is deployed in populated areas that are well served by the Internet.
- Internet communication is **always available**, **continuous**, and virtually **instantaneous**. It can readily convey telemetry immediately from sensing instruments to research centers.

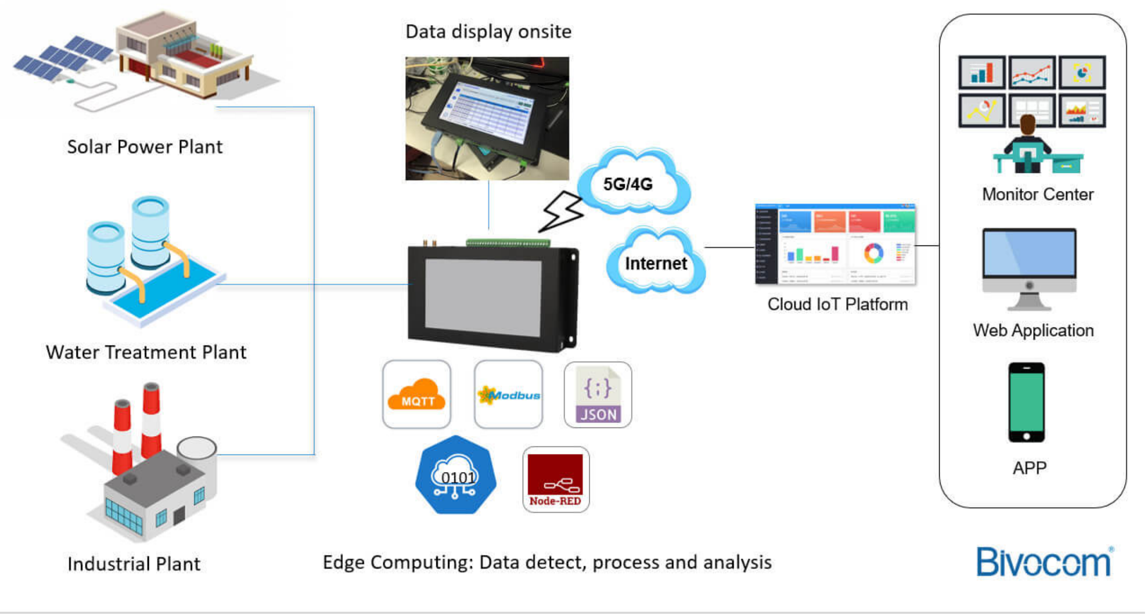

Remote Instrumentation

- But over 2 billion people lack access to the Internet, and the parts of the Earth that are uninhabited – including the oceans, which cover 71% of the planet’s surface – have no wired Internet infrastructure at all. Telemetry from environmental instrumentation where wired Internet is absent must be conveyed by wireless means.

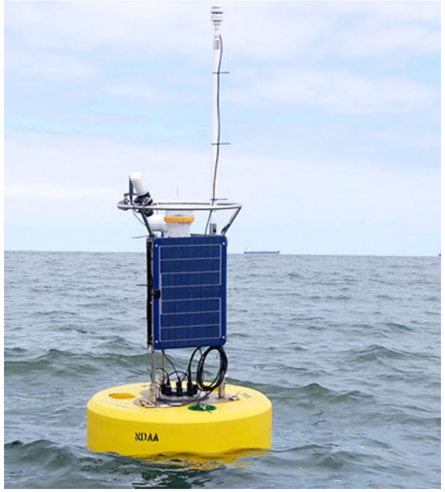

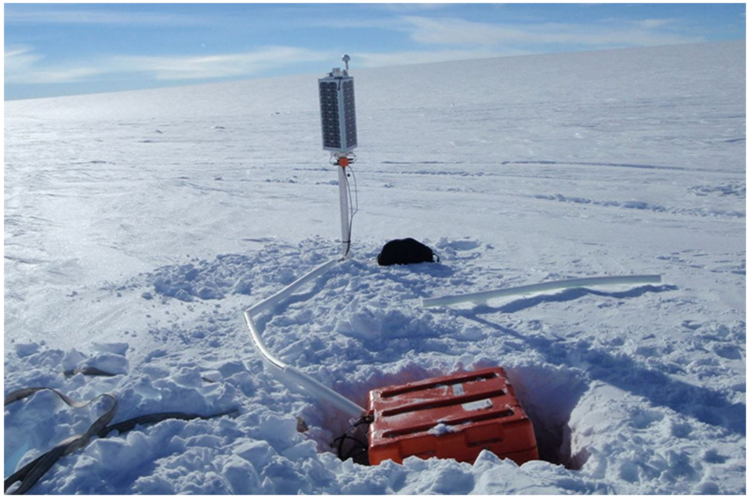

Constellations

- A large constellations of communication satellites in low-Earth orbit – for example, Starlink – can offer wireless Internet service to all points of Earth's surface within a wide band of latitude.
- But such service may be expensive, and it's not available in the uninhabited regions outside of that latitude band. Without paid subscribers, it's not commercially viable.

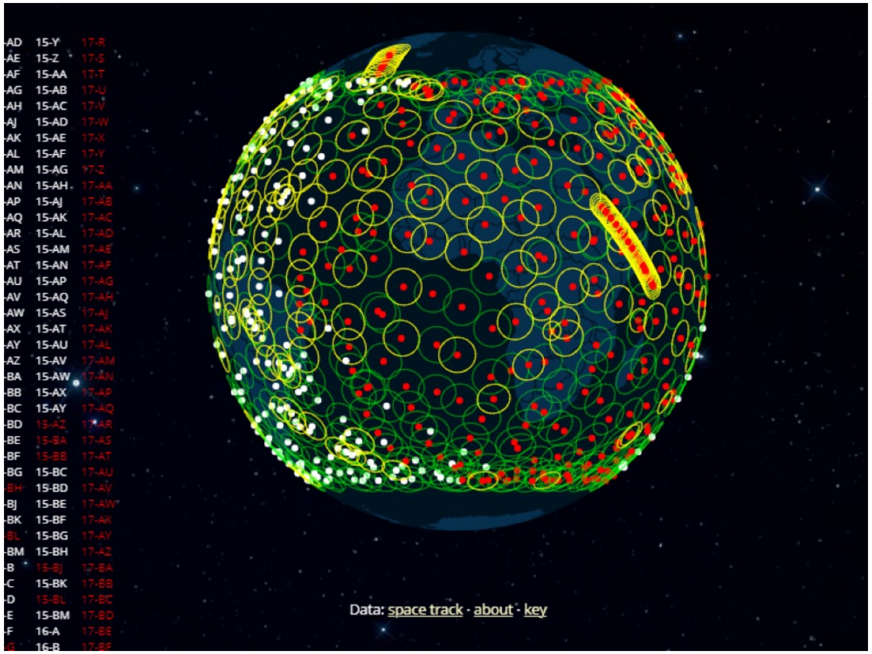

What do we do instead?

- There is an alternative: delay-tolerant networking (DTN) is network communication that is likewise **always available** – but **not necessarily continuous** and **not necessarily instantaneous**. Relaxing these service expectations enables DTN to provide *automatic, reliable, secure communications* without requiring wired infrastructure or access to satellite mega-constellations, therefore at lower cost.

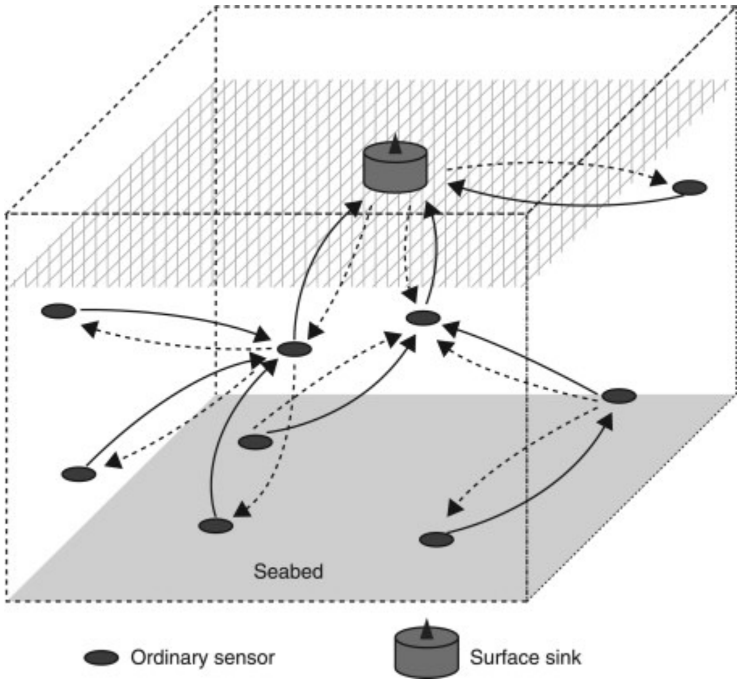

Origins of DTN

- DTN was originally developed to enable wireless network communications in deep space flight missions, where conditions would not favor the Internet architecture.
- Transmission wouldn't be continuous:
  - Orbital movement and planet rotation would sometimes make the reception of transmitted signals impossible.
  - Signal reception requires that the signal source and destination be accurately pointed at each other, but mission operations would make this impossible at some times.
- Transmission wouldn't be instantaneous:
  - No signal can travel faster than 299,792.458 km/s.
  - Mean distance from Earth to Mars is 225 million km.
  - So mean latency in reception at Earth of a signal issued at Mars is 16.3 minutes.
- A different network architecture was needed. Here's why:

It's All About Delay

- Network discontinuity is, in essence, highly variable delay.
  - Case 1: continuous connectivity but client is, e.g., 56 million miles from server. *Response to query always arrives 10 min. after query is issued*.
  - Case 2: client and server are in adjacent offices but router is powered off for, e.g., 10 minutes. *Response to query, on this occasion, arrives 10 min. after query is issued*.
- Key effect of delay: reliable transmission of a given byte of data can take an arbitrarily long time.
  - Transmission can be lost due to corruption, N times.
  - NAK can be lost due to corruption, N times.
  - Disruption can delay transmission of NAK (or retransmission of data) by an arbitrarily long time.
- So round-trip times – e.g, query/response – can take an arbitrarily long time.

Effects of Long and/or Variable Delay

- Connection establishment could take more time than entire communication opportunity.
  - So protocols must be connectionless.
- Transmission history can't be used to predict round-trip times.
  - So communication timeout interval computation must rely on link state information rather than timing statistics.
- End-to-end retransmission would reserve resources (retransmission buffer) at the originator and delay delivery of data to the destination for the entire duration of the transaction – possibly days or weeks.
  - So retransmission should be between relay points within the network rather than end-to-end: custody transfer.

Effects of Delay (cont'd)

- In-order stream delivery could be stuck for a long time, waiting for byte N to arrive before delivering byte N + 1.
  - So out-of-transmission-order delivery is needed – multiple concurrent transmissions.
  - So data must be structured in transmission blocks (e.g., files or messages) for concurrent retransmission – *not streams*.
- But reliable transmission of any single block can take an arbitrarily long time.
  - So any number of message transmissions might be in progress at the moment a computer is rebooted or power cycled.
  - So retransmission buffers should reside in non-volatile storage – not memory – to minimize risk of massive transmission failure.

Effects of Delay (cont'd)

- Propagation of changes in topology could take an arbitrarily long time.
  - So route selection that is based on knowledge of the current topology of remote sectors of the network is more likely to be wrong than in a low-delay environment.
  - So either:
    - Changes in topology must be limited, or
    - Routing must be based on some other information that is innately stable, such as endpoint <u>name</u> rather than last known address.
  - Limiting change in network configuration limits the usability of the network. It might work for a decade or so, but that would just be pushing the problem off until after we retire.
  - The long-term solution is late binding of endpoint names to network topology.

End-to-end IP in the Space Network

- For *unreliable* transmission, there are three problems:
  - Non-standard implementations would be needed to handle disconnection.
    - Network partition must not be treated as an enduring change in topology.
    - Network partition must cause outbound packets to be queued for future transmission rather than discarded.
  - The standard internal routing protocols use history-based timeouts to detect route failures. Routine transient partitioning would incorrectly cause route failure to be inferred and propagated, resulting in routing table errors.
  - Routing is based on IP address and known network topology, which might change too rapidly for routing protocols to track. Result would be routing errors.

End-to-end IP in the Space Network (2)

- For *reliable* transmission, there are additional problems:
  - Options are either (a) TCP or QUIC or (b) UDP with reliability implemented in each application or middleware.
  - TCP isn't suitable.
    - It's based on connections, streaming, end-to-end retransmission, in-order delivery.
    - Retransmission buffers are in memory.
    - Timeout intervals are computed from transmission history.
    - Data loss is assumed to be due to congestion rather than corruption, so performance over noisy wireless links is degraded.
  - So the BGP external routing protocol standard, which uses TCP, is not suitable. So none of the standard routing protocols work.

End-to-end IP in the Space Network (3)

- More *reliable* transmission problems:
  - QUIC isn't suitable. Retransmission timeout intervals can be tuned, multiple streams can run concurrently over a single connection, but:
    - It's still based on connections, streaming, end-to-end retransmission, in-order delivery.
      - Connections can be established at launch time and sustained indefinitely, but insertion of new network nodes after launch will still require new connections.
    - Retransmission buffers are still in memory.
    - Timeout intervals are still computed from transmission history.
    - Data loss is still assumed to be due to congestion rather than corruption, so performance over noisy wireless links is degraded.

End-to-end IP in the Space Network (4)

- More *reliable* transmission problems:
  - Without standard infrastructure for retransmission:
    - Each reliable application-layer protocol must reinvent retransmission – **additional cost and risk**.
    - **No standard for congestion control**.
  - Reliability from retransmission by the application is (like TCP) based on end-to-end retransmission. So:
    - **No custody transfer**. Network performance is degraded.
    - Timeout-triggered retransmission is excluded: end-to-end round trip time is not accurately predicted. **Distributed Consensus problem**.

Ruling Out the Internet Architecture

- TCP isn't suitable, for a host of reasons. Neither is QUIC.
- There's no alternative Internet standard for reliable transmission that would work over deep space links.
- So no standards for flow control and congestion control.
- None of the standard routing protocols would work.
  - BGP relies on TCP. Others rely on timers that won't work right.
  - Transient network partitioning mustn't be interpreted as topology changes.
- No COTS routers would work.
  - Interruption of outbound link must cause outbound traffic to be queued rather than discarded.
- All that's left is UDP/IP with static routing: just a less bit-efficient packaging alternative to raw CCSDS packets.

How DTN Differs from Internet

- Like Internet, DTN is an overlay network.
  - DTN “bundle protocol” (BP) is to IP as IP is to Ethernet.
  - A TCP connection within an IP-based network may be one “link” of a DTN end-to-end data path; a deep-space R/F transmission may be another.
- But reliability is achieved by retransmission between relay points within the network, not by end-to-end retransmission.
- Also, route computation has temporal as well as topological elements, e.g., a schedule of planned contacts. There is no expectation that the topology of the network is the same from one second to the next.
- The contact plan enables rate control instead of flow control, and congestion is predicted and averted rather than detected and repaired.
- And forwarding at a router is automatic but not necessarily immediate: store-and-forward rather than “bent pipe”.

DTN Stack Elements

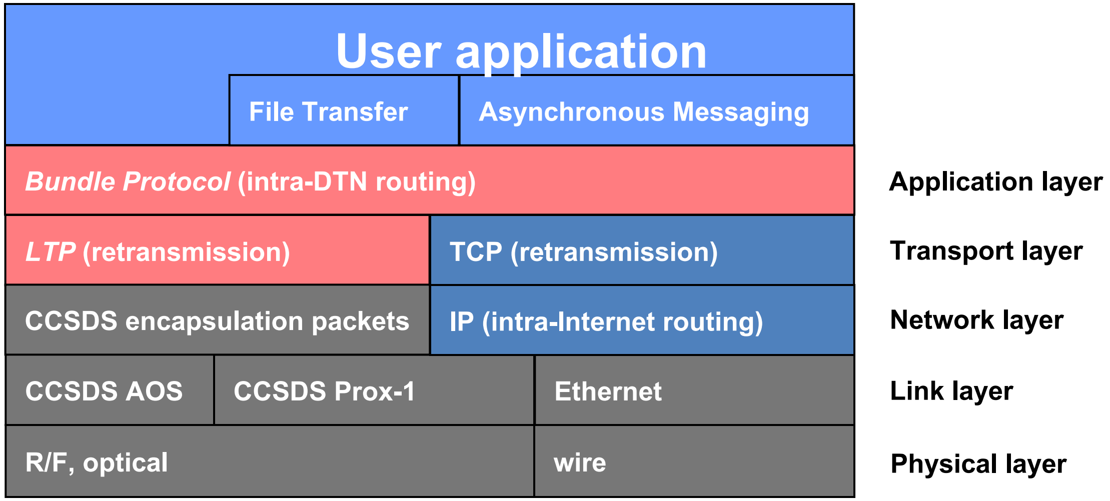

An End-to-end Path

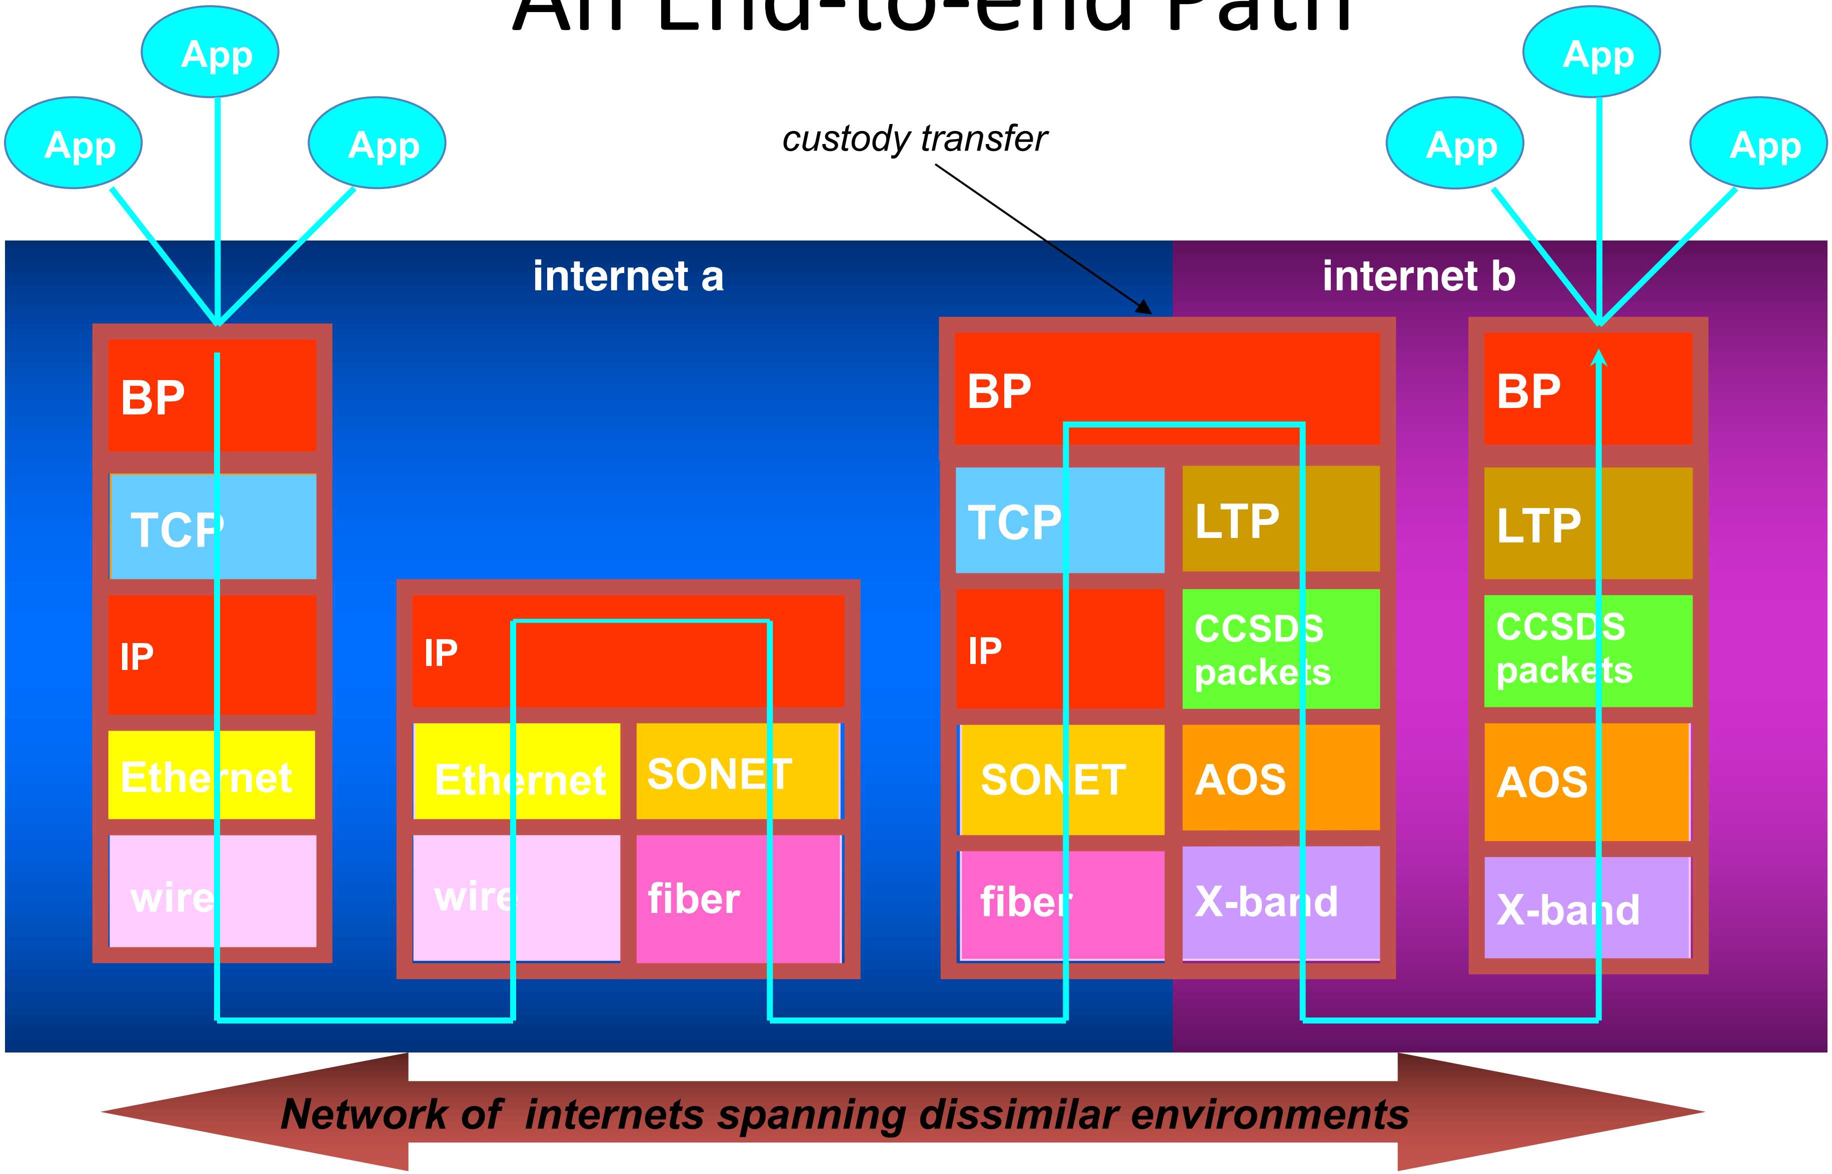

DTN for Challenged Terrestrial Networks

- As in deep space, some terrestrial communication regimes – for example, environmental monitoring – may not favor the Internet architecture.
- Transmission won't be continuous:
  - Some sensor nodes will not be in R/F range of one another. Transmission may be possible only when relay is enabled – a drone approaches or a satellite is overhead.
  - Sensor nodes may have limited power, preventing radios from running at all times. Recharging from solar panels may be inhibited by prolonged overcast.
- Transmission won't be instantaneous:
  - Transmission to a relay will be immediate, but transmission from the relay vehicle to the research station can occur only when it comes into R/F range. This may not occur until after the relay vehicle has physically traveled some number of kilometers.
- A different network architecture is needed.

Questions?