Internet Resilience

Geoff Huston AM

Chief Scientist, APNIC

How?

- Do we incorporate resilience into the Internet's infrastructure, protocols and applications?
- What are the practical limits on engineering resilience?

Transmission Reslience

How do you protect cable systems?

- Use a mesh of internal connectivity such that each node is connected to 2 or more other nodes
- Use a control system that can detect inter-node link failures and re-direct data flows around the failure

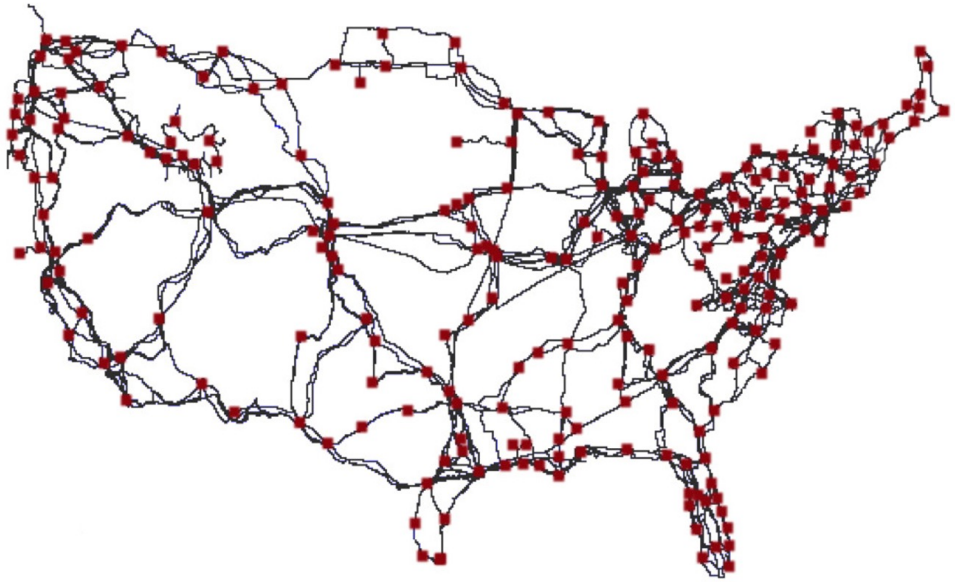

https://pages.cs.wisc.edu/~pb/tubes_final.pdf

Transmission Resilience

- How do you protect undersea cable systems?
  - By using a large collection of cables so that no single cable becomes a single point of vulnerability

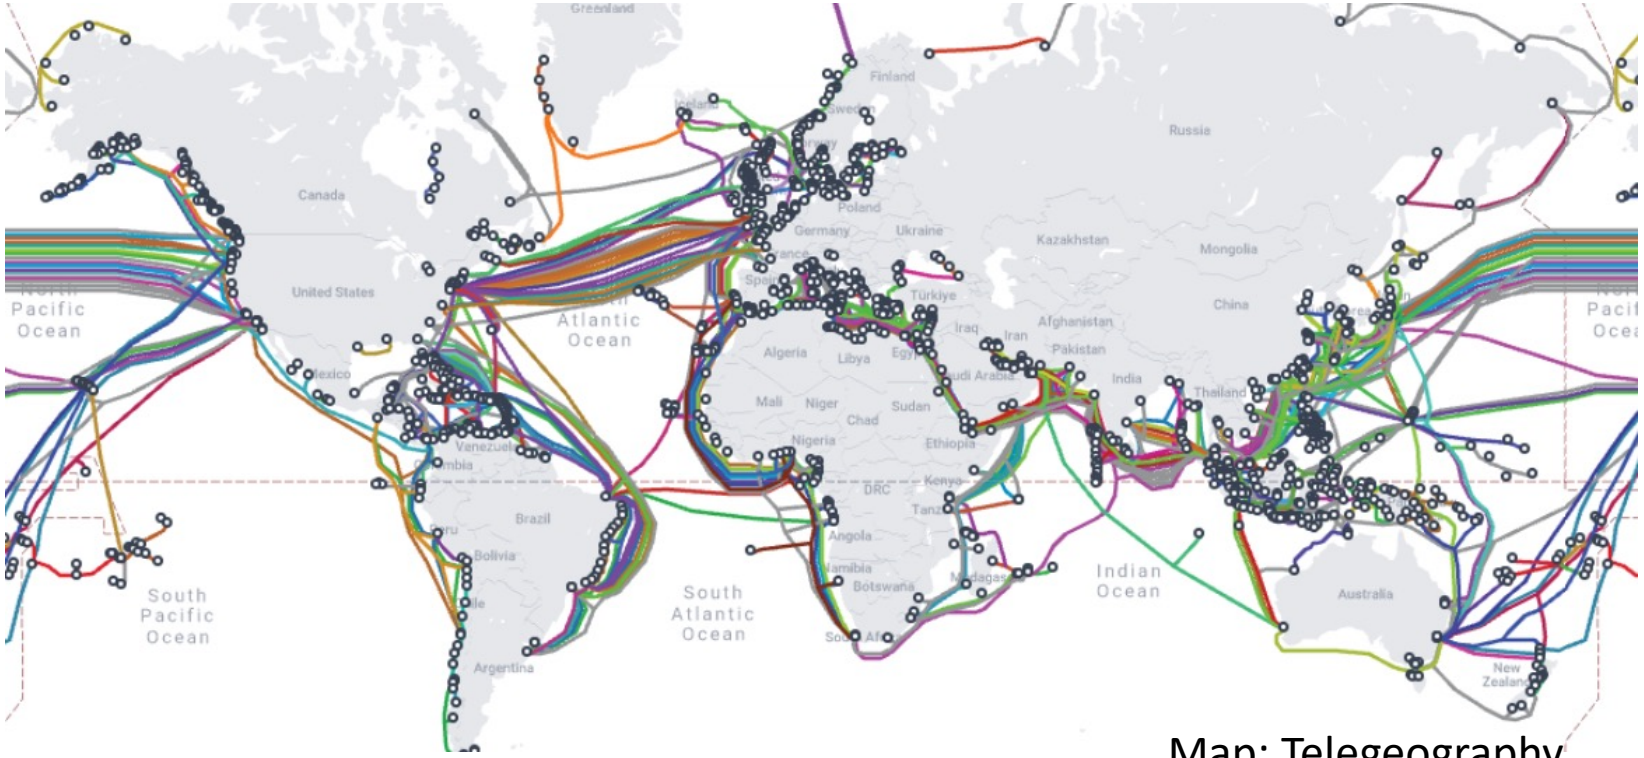

Map: Telegeography

Transmission Resilience

- Low Earth Orbiting services
  - These compound systems can perform mutual backup across the satellite constellation – no single satellite is critical to the service

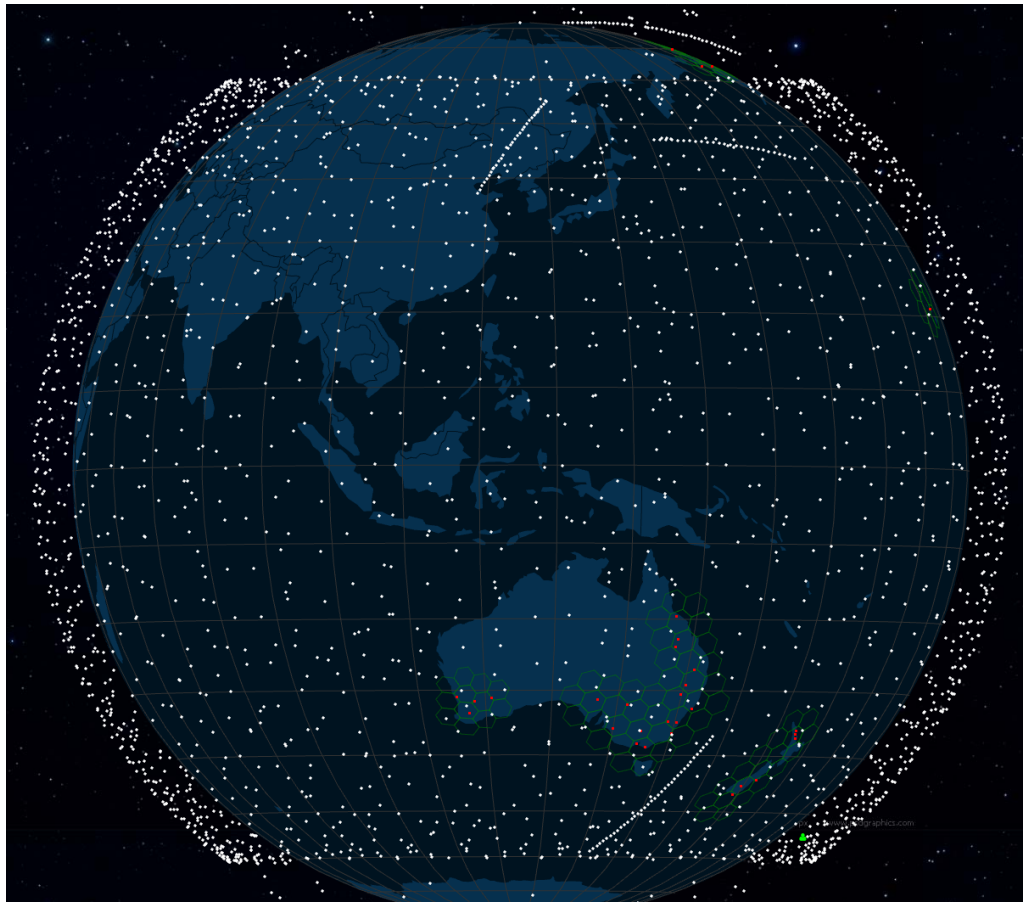

satellitemap.space

Transmission Resilience

- MEOs and GEOs
  - Higher altitude spacecraft do not operate in a dense mesh so they generally do not offer mutual backup
  - They can complement terrestrial systems by offering a space-based path as a backup

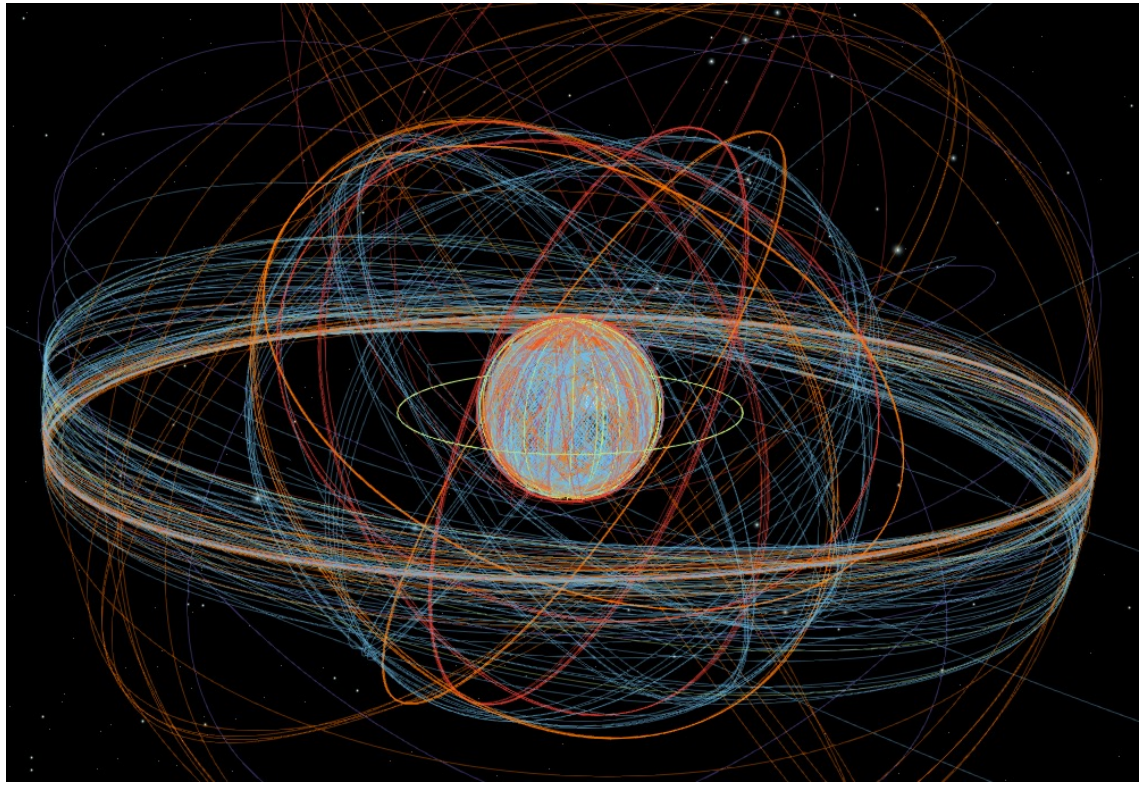

https://geoxc-apps.bd.esri.com/space/satellite-explorer/

How to support transmission resilience

- It you have a self-repairing transmission system, then the issue becomes how to redirect data flows along new paths when failure occurs
- Stateful virtual circuit systems have a more challenging task here as they have to reload a coherent set of forwarding directives to restore virtual circuit integrity
- Stateless packet-based forwarding systems can respond more directly to dynamic changes in the network’s internal topology

Datagram Packet Networking

- By allowing each node greater ability to determine how to handle individual communications elements packet-based networks can be used in systems that are resilient against many forms of node and link failure
  - as long as the basic connectivity infrastructure is sufficiently dense
- Constructing a network using a packet-based foundation, is possible to construct a highly resilient and reliable system from less reliable components

Packet Networking

- Divide a communications stream into a sequence of individual “packets”
- Add the intended destination as an addresses to each packet’s header, together with a sequence identifier and pass the packet into the network
- Allow the destination to reassemble the communications stream from the sequencing information in the received packets
- Each packet in an independent transaction, and packet loss can be repaired individually
- The network has no end-to-end virtual state to support the packet stream – each packet is handled independently
- Packet loss can be tolerated without needing to shutdown and restart the entire transaction

Network Routing

- The second part of network resilience lies in the routing system
- Routing uses a distributed computation for topology maintenance
- Within each node the topology is represented as a set of all reachable destinations and an associated net-hop forwarding decision
- There is no single control point and no single point of failure
- The routing system detects failures and mends the connectivity fabric

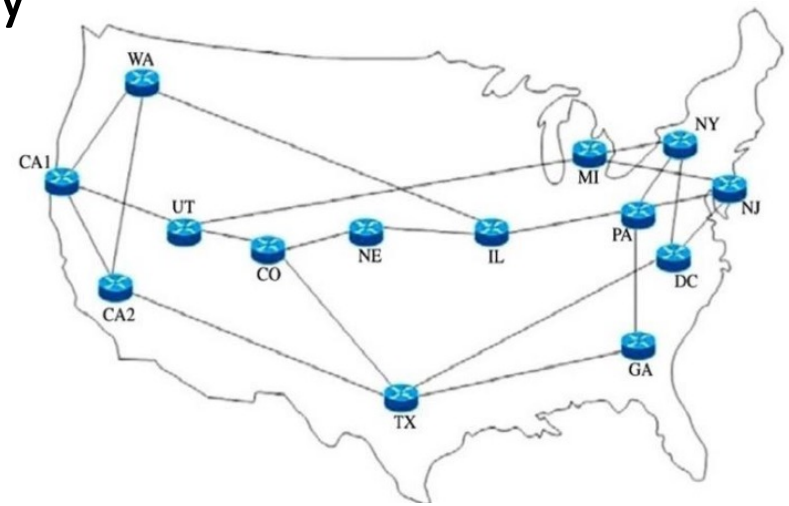

NSFNET c. 1993

IP Resilience

- How can you support resilience at the edge of the network as a customer?
  - Use multiple service providers
  - Announce the site's addresses to each provider
  - Import the full reachable route set from each provider

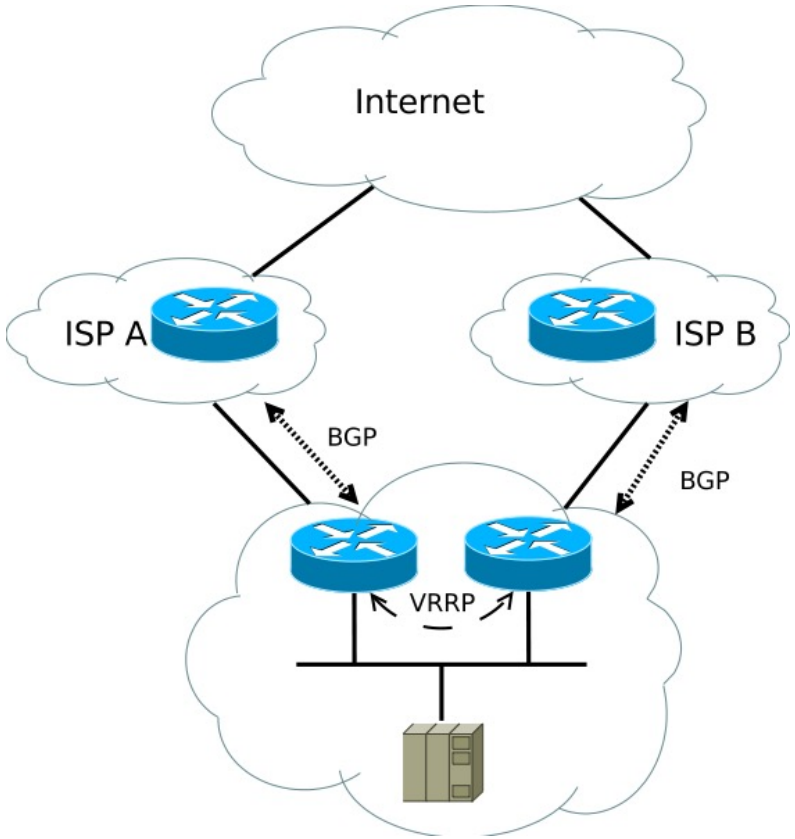

IP Resilience

- In IPv6 we thought that a multi-homed site would use address prefixes from each upstream ISP and each internal host would have multiple IPv6 addresses
  - But this presents a challenge to support seamless handover in the event of failure as active sessions need to renumber
- SHIM6 was an effort in IPv6 to allow a host to exploit a network's multiple external connections and implement transparent failover across external paths

Transport Resilience

- TCP incorporates loss recovery and rate adaptation into the basic protocol operation
- MultiPath TCP allows a host to use multiple network interfaces to load share across multiple paths – it can also support some forms of failover
- QUIC allows for address agility, where a session can adjust to external network events that trigger a change in IP address
  - NAT session re-binding
  - Host address rebinding
- HIP used similar concepts to binding a session to a shared session token, allowing the host addresses to change dynamically

Application Resilience

- DNS resilience:
  - Multiple name servers
  - Multiple recursive resolvers
  - Multiple IP addresses for address resource records
  - Keep querying until you get an answer
- Service resilience
  - Service replication through anycast networks
    - Anycast allows the same service IP address to be injected into the routing system at multiple points. Each client is directed to the “closest” insrtance of the service through the normal actions of the routing system
    - Individual instances can be added and removed with minimal disruption

Engineering Resilience

- The key feature of resilient engineering is to avoid single points of critical failure
- This requires:
  - over-provisioning of the service infrastructure to allow failover
  - flexibility in the service overlay to allow dynamic reconfiguration
  - Feedback-based control systems to allow for failure detection and healing

Engineering Resilience

- We have built approaches to support resilience at the transmission level, IP level, the transport level, the session level and the application level
  - Is this a case of too much? Do all of these responses trip over each other?
  - Or are they each responding to different classes of potential failure modes?
  - Are these systems too slow to react?

    SDH promised 50ms reaction times to primary path failure. BGP manages some 50 seconds on average to converge on a new path. The DNS may take up to 8 seconds to find a responding service.
  - What are the performance targets for resilience?
  - What are the costs of resilience?
  - Who pays?

Engineering Resilience

- Is is possible to position resilience as a competitive differentiator between ISPs? Or between products? Or applications?
  - Yes.
  - And no.

Internet Resilience

- The internet can support highly resilient services, but there are a few preconditions to allow it to work as intended:
  - Do not build single points of failure into your design
  - Provision infrastructure services with redundant capacity
  - Avoid stateful services

Thank You