TCP's Reaction to Soft Errors

Fernando Gont

(presented by Pekka Savola)

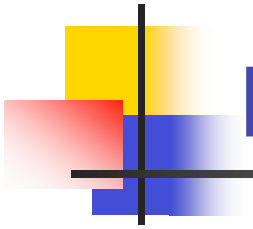

Error handling in computer networks

- Fault Isolation

  Actions to determine there's a network error.

  Performed in the Internet Architecture by ICMP.

- Fault Recovery

  Actions to survive a network failure (if possible).

  May depend on the type of error being reported, the time the error is reported, and any extra knowledge about the most likely fault scenarios.

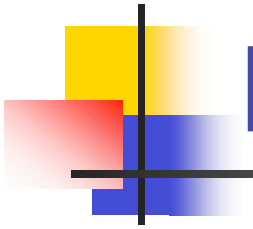

Error handling in TCP

Static fault-recovery policy

- Hard Errors

  Abort the connection

- Soft Errors

  Record the error, and repeatedly retransmit the segment until a retransmission succeeds or the connection times out

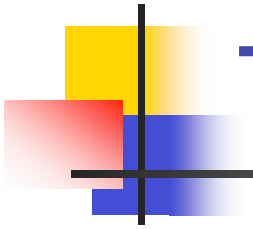

Things that may go wrong

- ✓ A domain name maps to several IP addresses
- ✓ Some of them are unreacheable
- ✓ TCP's policy says we must retransmit until the connection times out
- This would imply long delays between connection-establishment attempts

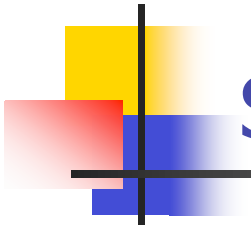

Scenarios where things may go wrong

- IPv4

  Packets may be black-holed (no notification)

  Destination unreacheable (notified by ICMP)

- Dual Stack IPv6 on by Default

  Packets may be black-holed (no notification)

  No default route, and NUD fails (notified by ICMP)

  Scope boundary enforced by firewall (notified by ICMP)

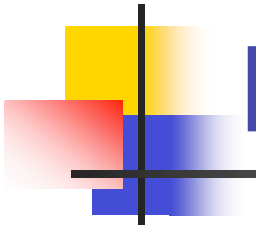

Possible solutions

- Change TCP's reaction to soft errors in the SYN-SENT and SYN-RECEIVED states
- Asynchronous application notification
- Provide applications a Higher-Level API

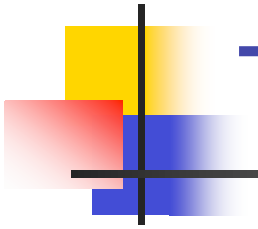

Things we must consider

- **Who** should decide which is the best solution?
  - Protocol Developers?
  - System Administrators?
  - Application Programmers?
- **When** do we want the fix to begin to be used?

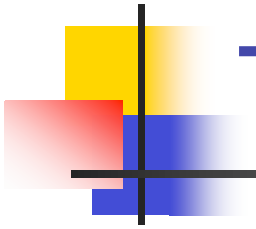

Things we must consider

- What do **robust** applications do?
  - Interactive applications

    Is there any ***real*** success in connection establishment after several minutes?

  - **Non-interactive applications**

    Don't they have to implement retry mechanisms, anyway?

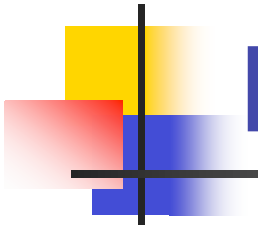

Next Steps

- Adjust other similar protocols (SCTP, DCCP?) as necessary
- Take as WG document, for BCP (or Informational)?