

Part 2

The Future

New insights

- A better way to start
 - Start in “verification phase”
 - Best guess MTU as “provisional MTU”
 - A safe fallback MTU
- Another view of layering and shared state
 - Not fully sorted out
- New combined transition and verification phase
 - Can overlap probing and verification
 - Greatly speeds convergence

Other possible ICMP changes

- Long ICMP attacks thread on TCPM
 - ID Gont-tcpm-icmp-attacks
 - Describes methods to make PTB processing more robust
 - Add TCP sequence and ACK number checks
 - Add suggested TCP sequence checks into pmtud
 - Add language about other optional checks
 - IP ID or IP image checks

The next steps

- Fully catch up on NITs and other input
- More progress on the implementation
- Fold in insights
- Preliminary MIB
- Advocate other implementers

Algorithm Review

(Updated slides from prior meetings)

Packetization Layer path MTU discovery

- *New* Start in verification phase
 - mtu=best guess from route or interface
 - Search_low (fallback)=1kB
- Probe with larger packets to test MTUs
 - Provisionally raise MTU if successful
 - (Optional) Process RFC1191 ICMP
 - Do not reduce TCP window on lost probe
- Verify provisional MTU for some interval
 - Additional losses imply problems

Layered Implementation

- State kept in path information cache in IP layer
 - Probing state and timers
 - Recent successful and unsuccessful probe sizes
- Algorithm runs in the Packetization Layer
 - PL cuts the data into packets
 - Probing and verification are intrinsically PL specific
 - New description facilitates sharing the rest of the code
 - The search heuristic and error logic can be shared

Key Properties

- Robust
 - Tolerates ICMP delivery problems
 - Verification phase addresses spurious delivery
- Progressive interoperation with classical pMTUd
 - Start large and process all ICMP
 - Start small and ignore all ICMP
- Parallel to congestion control
 - End to end algorithm: use loss as the feedback to adjust window or packet size
 - Well understood limitations

Robust

- Primary design goal: Do no harm
- Avoid problems with RFC 1191 pMTUd
 - Not affected by ICMP delivery problems
 - Not affected by tunnels and encapsulation
 - Not exposed to RFC 2923 problems
- Minimal new exposure
 - Spurious delivery of oversized packets
 - Verification phase provides protection

Progressive deployment

- Enhance RFC1191 pMTUd
 - Start with large MTU and process ICMP
 - Use PLPMTUD iff repeated timeouts
 - Maximally robust from a deployment perspective
- Replace RFC1191 pMTUd
 - Start with small MTU, ignore all ICMP PTB messages
 - Search upwards to raise MTU
 - Maximally robust from a security perspective

Parallel to Congestion Control

- End-to-end algorithm
- Adjust data stream parameters:
 - Packet or window size
- Use packet loss for feedback
 - Interactions with Congestion Control are specified in RFC2119 standards language
- Better fit with end-to-end principle(?)