


# Challenges deploying PTPv2 in a Global financial company


[http://tagus.inesc-id.pt/~pestrela/timip/Challenges\\_deploying\\_PTPv2\\_in\\_a\\_Global\\_Financial\\_company.pdf](http://tagus.inesc-id.pt/~pestrela/timip/Challenges_deploying_PTPv2_in_a_Global_Financial_company.pdf)

Pedro V. Estrela, PhD  
Performance Engineer  
05-Nov-2013



# Financial industry landscape

- **IMC: Global Liquidity provider**
  - Market-Maker in Equities / Derivatives markets
  - Provide best prices to the other participants
  - 4 offices worldwide
- **Time Sync needs**
  - Regulation Compliance
  - Risk Management
  - Performance Optimization




# Initial global NTP situation

- Financial centers

- Major locations: Chicago, Frankfurt, London, New York, Seoul, Sidney

- IMC private network:

- All global timezones
- 10s of datacenters + leased lines
- 100s of legacy + PTP-aware NEs
- 1000s of end trading-servers



# Basic WAN support

- Multicast TTL>1:
  - One GM for all timezones !
- Clock separation options:
  - a) Fine tuning TTLs
  - b) PTPv2 Sub-domain field
  - c) Traffic blocking (ACLs)
  - d) Separate v1 groups

# Scalability issues

- Multicast profile

- Upstream overhead


- (1000s of clients, with few BCs or TCs on the network)

- Endless problems from “All-to-All” semantics

- Unicast profile

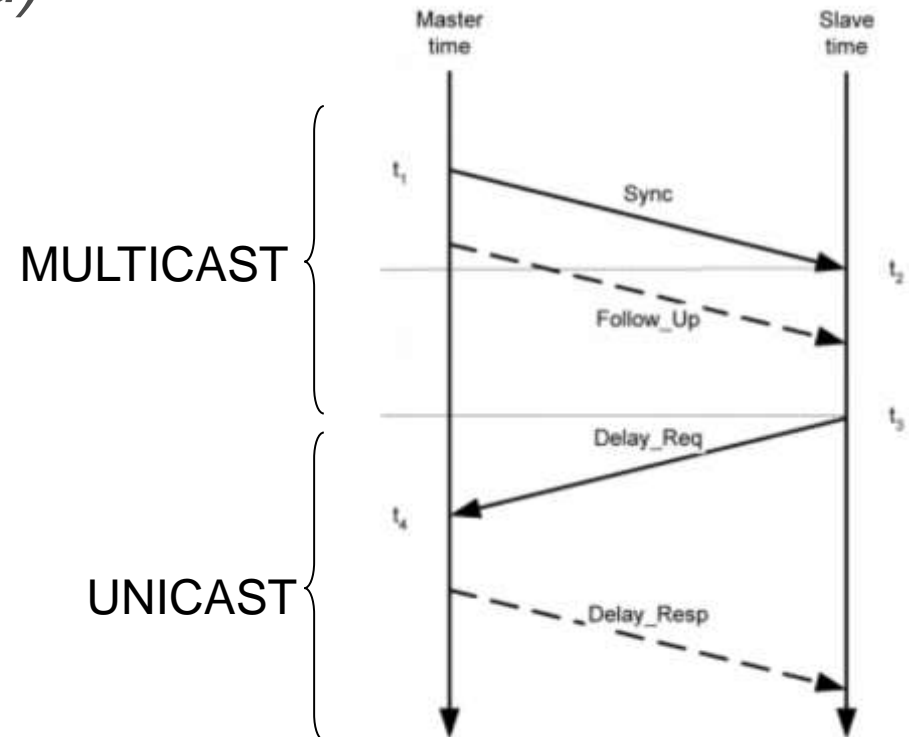
- Downstream overhead

- Poor support




# Hybrid mode solution

- Downstream: regular multicast
- Upstream: unicasted directly to the GM
  - *(already contributed to PTPd)*



# Robustness problems pt1

- GPS jamming
  - NPL UK time distribution
  - CME / NIST time distribution
- Leap seconds
  - Linux kernel bug
  - Recurrent abolishment proposals





# Robustness problems pt2


- GM “traitor” scenario:
  - GM sent bad time (leap seconds = 0)
  - Backup GMs stay passive (same BMC)
  - Clients trust their single GM = jumps / slews
- Complete solution:
  - Always corner cases with single GM
  - Clients must listen to  $2^*T+1$  sources (1997 proof)





# IETF Enterprise profile

- Replace NTP in large companies?
  - UTC to end-user applications
  - Smooth migration!
  - WAN Focus
- Requested:
  - Scalability -> Hybrid mode
  - Robustness -> Multiple time sources
  - Accuracy -> Slave WAN jitter filters



# Conclusions

- Issues on PTPv2 standard itself:
  - WAN Clock separation (single mcast group)
  - Scalability: Hybrid Mode
  - Robustness: GM is single point of failure
  - Accuracy: Slave jitter filters
- Leap seconds
  - Expect continuous problems here!

