



# STAC Update: Fast Data

Peter Nabicht  
President, STAC

[peter.nabicht@STACresearch.com](mailto:peter.nabicht@STACresearch.com)

# Overview

- FPGA Special Interest Group
- Cloud connectivity benchmark development
- STAC-N1 (full network stack)

# FPGA SIG Background

- Founded with 7 financial firms
- Goal:
  - Work together on non-proprietary challenges in FPGA development that all firms would benefit from solving
- Initial objectives
  - Facilitate dialog regarding common challenges in FGPA design, development, testing and deployment
  - Articulate industry requirements for FPGA hardware and toolchains where commonalities exist

# What's happening now

- The group is meeting every 4 to 6 weeks
- Grown to
  - 16 financial firms
    - Exchanges, hedge funds, prop shops, and banks
  - 7 vendors
    - board, chip, development tools, and IP providers
    - vendors who leverage FPGA for their own projects
- Expanding on initial objectives, with focuses on:
  - Joint initiatives that allow for collaboration across financial firms and vendors
  - Exploring open source and open-source friendly projects
  - Deeper dives with vendors on critical tool chain components

# Current collaborations: 3 main projects

- RapidWright / RapidStream improvements, including
  - Common requirements, requests, and prioritized bugs
  - Collaborating with developers at AMD at a deeper level
- Language support
  - Jointly contribute to VHDL and SystemVerilog projects that check canonical language feature support in other tools
  - Use to convey of critical features to vendors
- Joint development of open-source Switch and/or NIC reference implementation
  - Exploring currently existing projects as starting points
  - Focus on the primary needs of trading firms

# Education

- Previously
  - Financial firms FPGA developers presented different build, test, and deploy pipelines
  - RapidWright project deep dive led by project engineers from AMD
- Upcoming
  - Tutorial for CXL for FPGA to CPU communication and impact on development from Intel


# FPGA SIG update

- Topics and projects are driven by interests of financial firms
- You too can join us

[www.STACresearch.com/fpga](http://www.STACresearch.com/fpga)

# Cloud connectivity latencies

- Have had increasing interest in understanding:
  - Latency
  - Determinism




**On-prem  
to Cloud**




# Cloud connectivity latencies

- Have had increasing interest in understanding:
  - Latency
  - Determinism




**On-prem  
to Cloud**



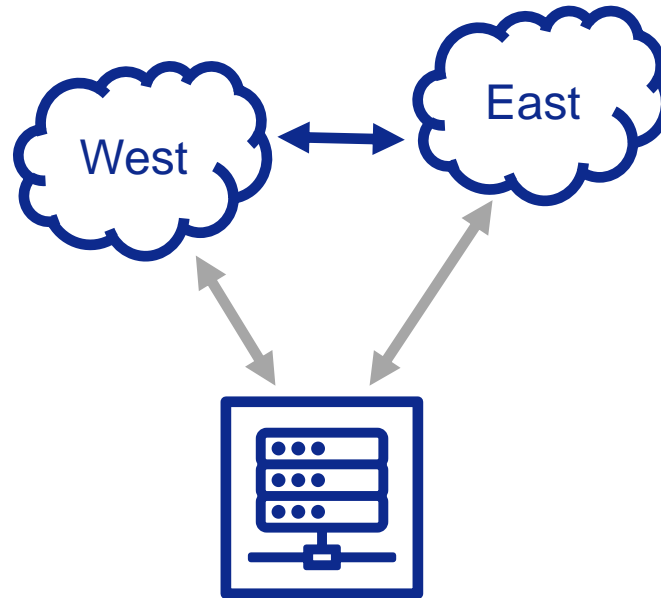
**Cloud Region  
to Cloud Region**

# Cloud connectivity latencies


- Have had increasing interest in understanding:
  - Latency
  - Determinism



**On-prem  
to Cloud**



**Cloud Region  
to Cloud Region**




**Cloud Provider  
to Cloud Provider**

# Issues with measuring cloud networks

- Opaque network infrastructures
- Dynamic network infrastructures
- Noisy neighbors
- No line capture
- Time synchronization?
  - What should be required for time synchronization?
  - What sub tests are needed to prove the accuracy of this?

# High-level diagram: on-prem to cloud

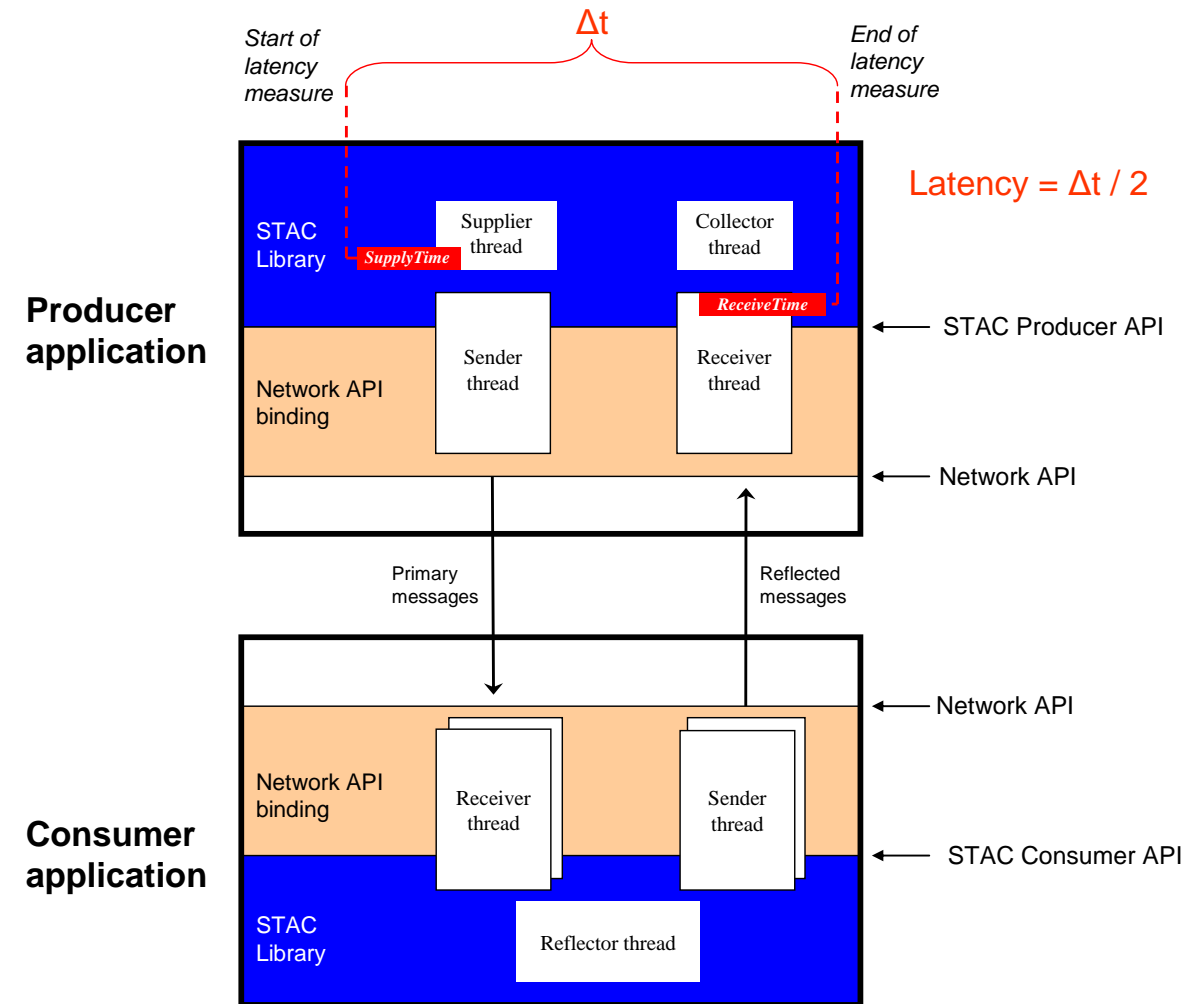


# Open questions on measuring cloud systems

- How long and when should we measure?
  - Time of day and time of month can impact performance due to noisy neighbors
- How many instances should we measure?
  - Different instances will have different paths to get to them
  - How many do we run in parallel vs how many total?
- Time synchronization?
  - What should be required for time synchronization?
  - What sub tests are needed to prove the accuracy of this?


# STAC-N1

- Measures the performance of a host network stack (server, OS, drivers, host adapter)
- Round-trip software timestamping
- Market data style workload
- Network API to network API
  - No middleware, feed handlers, etc.



# STAC-N1 / UDP / AMD / HPE / XtremeScale / OpenOnload


- First STAC software latency benchmarks with AMD EPYC
- Stack
  - STAC-N1 UDP-TCP binding
  - 2 x HPE ProLiant DL345 Gen10 Plus Servers
    - 1 x 32-core AMD EPYC™ 75F3 @ 2.95Ghz (4 GHz Boost)
    - AMD Xilinx XtremeScale™ X2522-25G-PLUS Adapter
    - Red Hat Enterprise Linux 8.4
  - 25Gb (via cross-over cable, FEC off)



[www.STACresearch.com/AMD221001](http://www.STACresearch.com/AMD221001)

# Vs. all public results for UDP on non-overclocked servers

- The lowest maximum latency for the base rate of 100k messages per second
  - [STAC.N1.β1.PINGPONG.LAT1](#)
- The highest maximum throughput tested of 1.2 million messages per second
  - [STAC.N1.β1.PINGPONG.TPUT1](#)
- The lowest 99p and max latency at the highest rate tested. Both:
  - [STAC.N1.β1.PINGPONG.LAT2](#)
  - [STAC.N1.β1.PINGPONG.LAT3](#)



[www.STACresearch.com/AMD221001](http://www.STACresearch.com/AMD221001)