Reflective Servers:
Seamless Offloading of Resource Intensive Data Delivery

Samuel Clark, Brookhaven National Laboratory

October 12th, 2023

Purpose

- Data updates from some 77,000 devices with >1M control points
- Many devices with *very* limited resources… how to do so efficiently?
- Proxy clients to devices through Reflective Servers
  - Offload asynchronous data delivery workload
  - Robust connection management
  - Easy deployment & configuration
  - Seamless integration into controls ecosystem

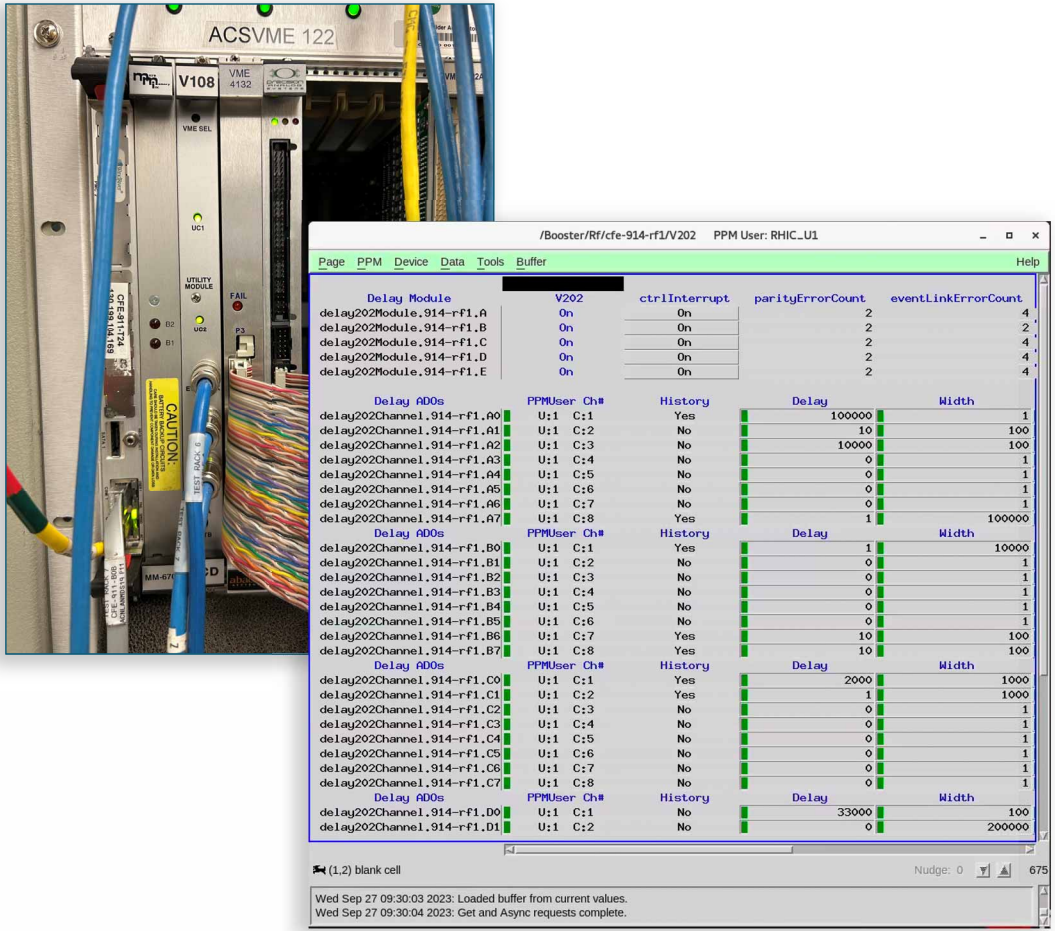

Brookhaven National Laboratory

Purpose

- Data updates from some 77,000 devices with >1M control points
- Many devices with *very* limited resources… how to do so efficiently?
- Proxy clients to devices through Reflective Servers
  - Offload asynchronous data delivery workload
  - Robust connection management
  - Easy deployment & configuration
  - Seamless integration into controls ecosystem

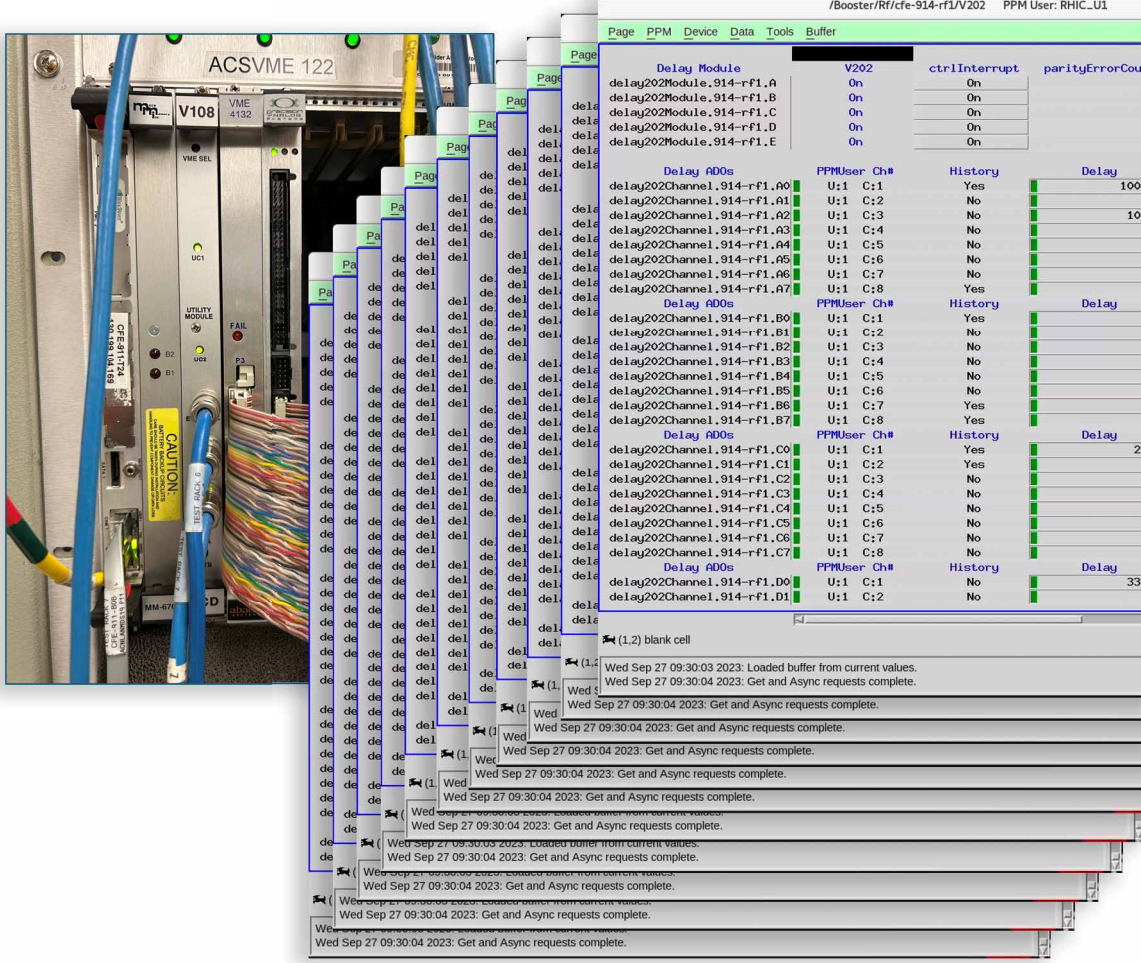

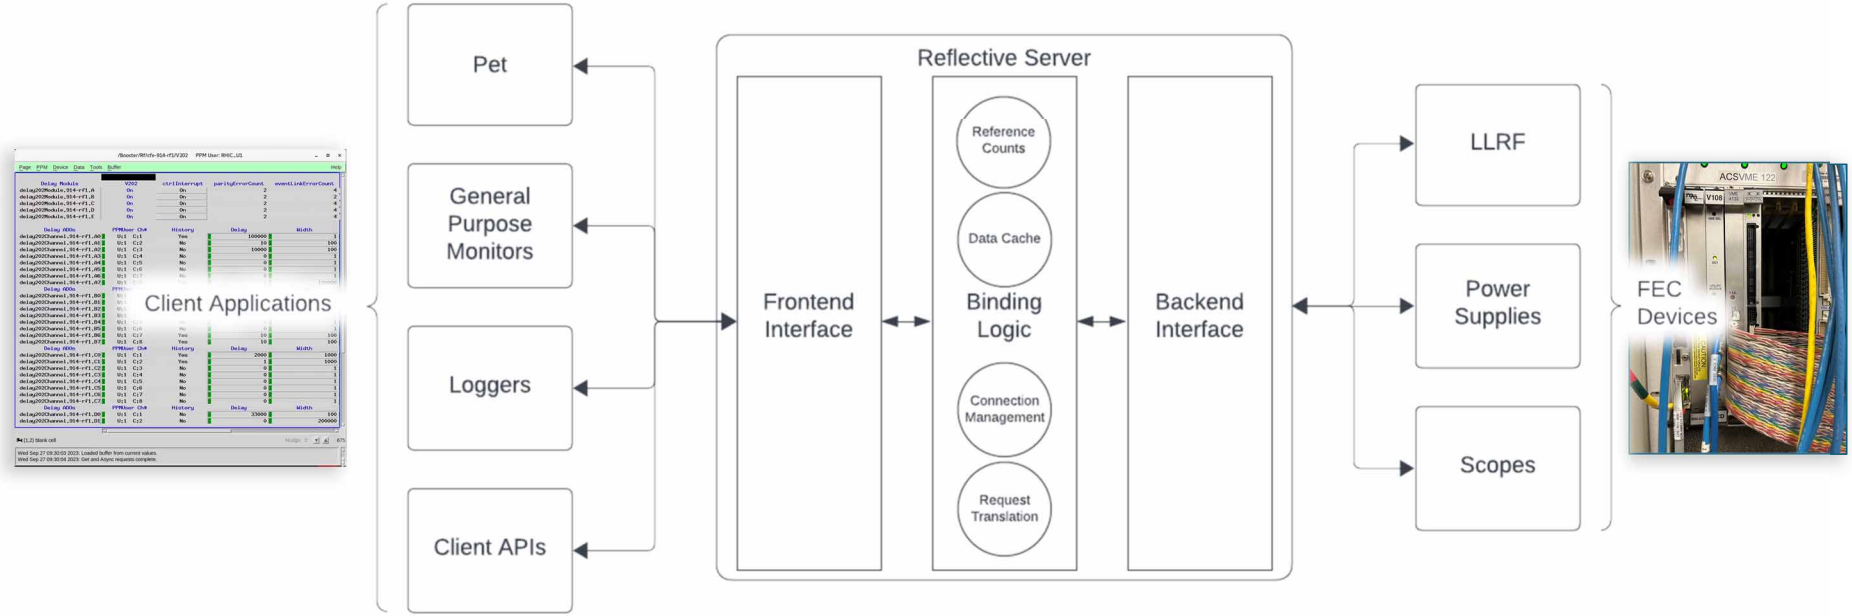

Implementation

- Reflective Server implemented in Python
- Fan-out data updates to 100s of clients
- All protocol features supported
- Efficient connection management
- Modular codebase for protocol flexibility
- Internal optimizations to reduce device traffic

Brookhaven National Laboratory

Performance

- Ability to handle much higher client load than FECs directly
- Analysis shows ~2x latency penalty for using Reflection
- Latency measures include…
  - Network round-trip time
  - Data marshalling in Reflective Server
  - Request/Response processing
  - Data fan-out in binding logic

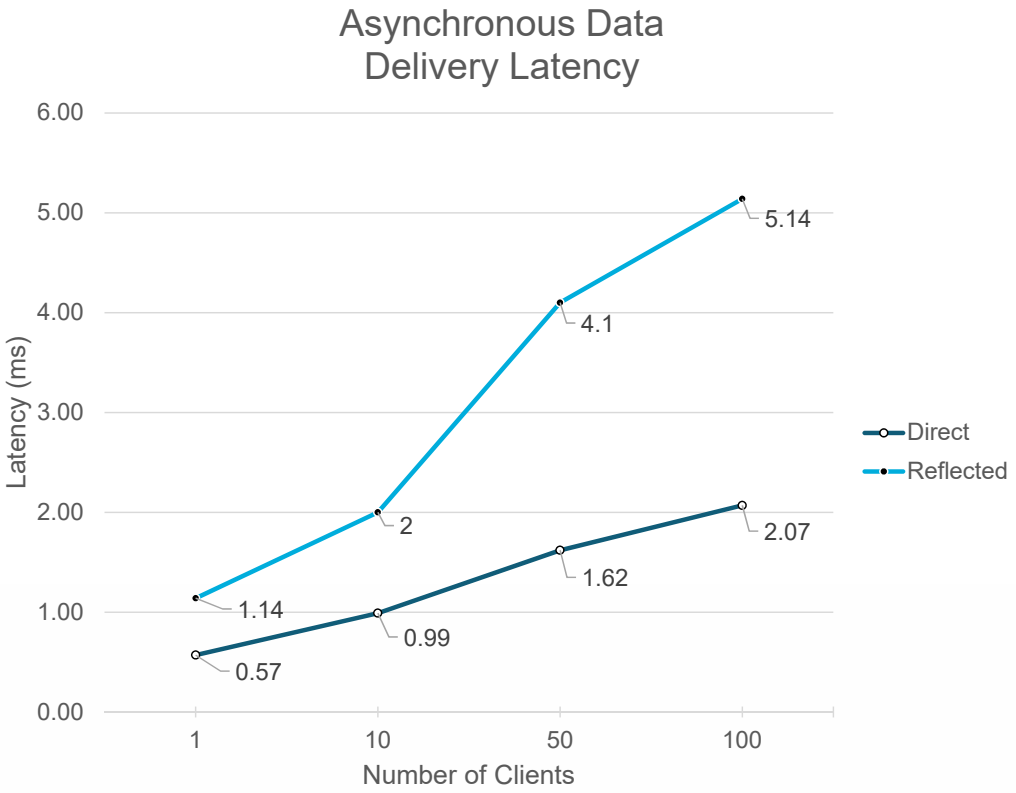

Brookhaven National Laboratory