



IP Load Balancing

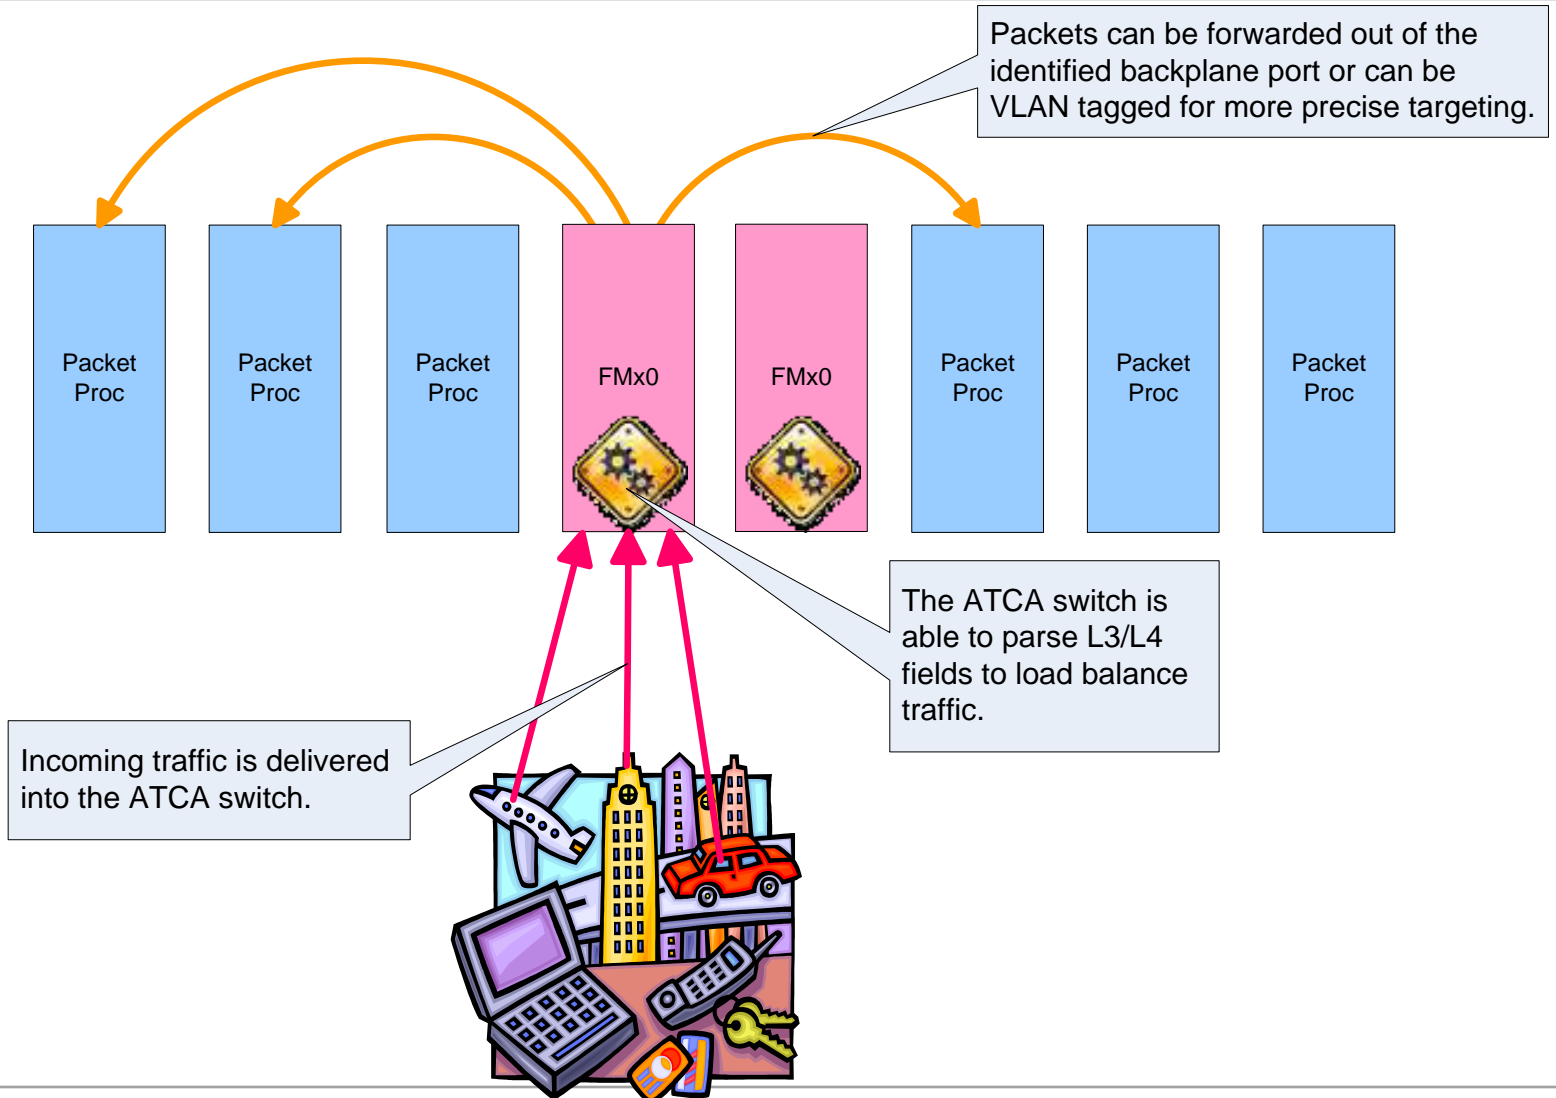
A critical part of system integration

James.Radley@ccpu.com

The Need for Load Balancing

- Deep Packet Inspection (DPI) applications are expected to handle bi-directional traffic on up to 4 x 10GbE links
 - “Bump-in-the-wire” traffic for shaping
 - Bi-directional wire taps for traffic profiling & legal intercept
- No one packet blade can handle all the traffic
 - Load must be shared across numerous packet blades
- Traffic from an individual user or flow (in either direction) will typically need to always pass through the same packet (DPI) blade
 - Load must be symmetrically handled

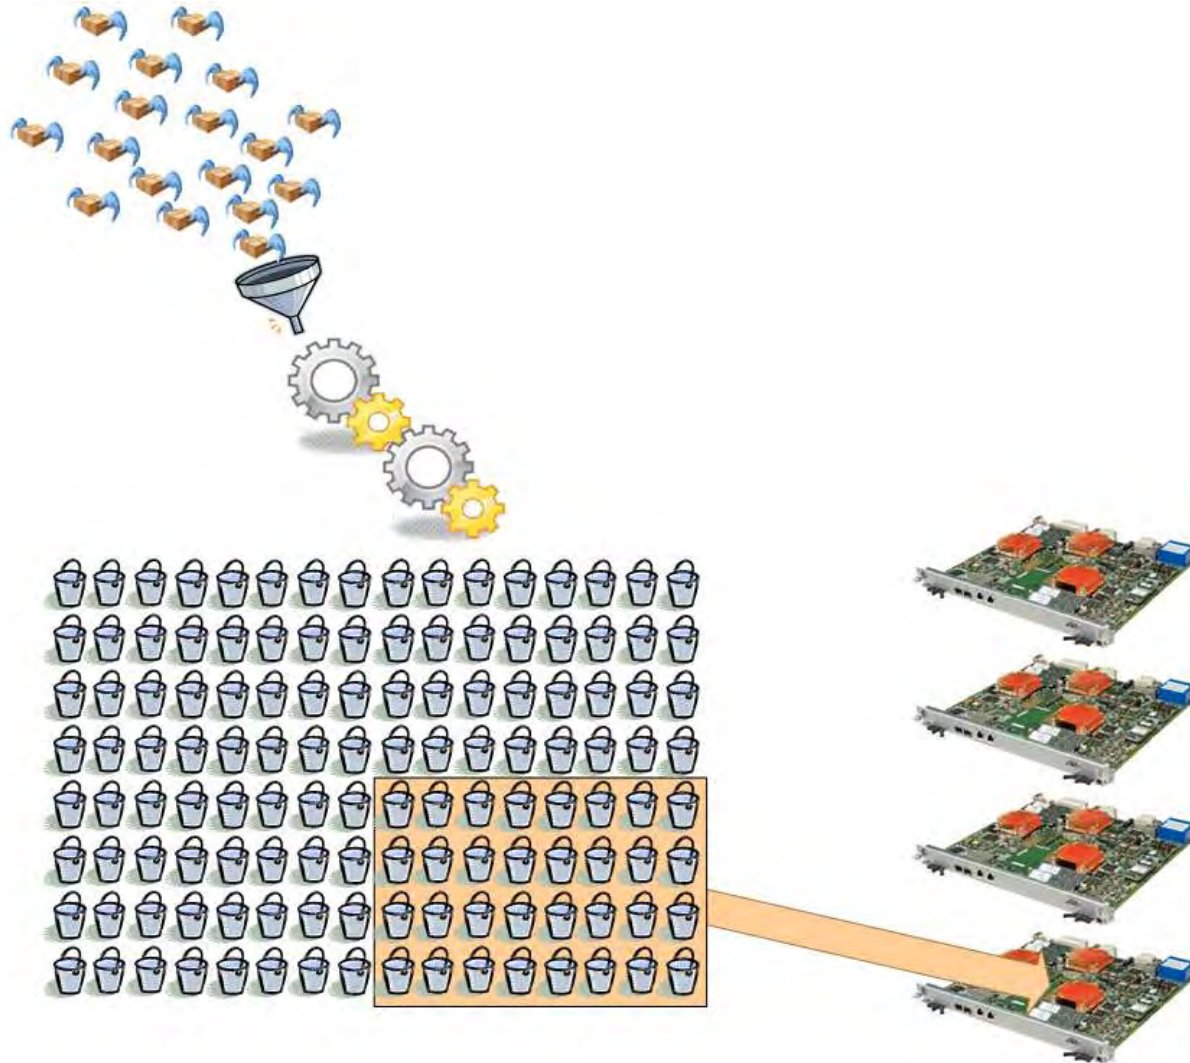
Load Balancing within an ATCA Service Platform



Service Platform: Critical Attributes for LB'ing

- Service reliability
 - Additional capacity available within the platform can be rapidly switched in to handle traffic displaced from a failed service blade
 - N+1 → Active / Standby
 - Spare blade ready to take over if another fails
 - N+M → Active / Active
 - Spare capacity maintained on all blades can be called upon to share load if a service blade fails
 - Failover must not affect existing traffic affinity across surviving blades
- Service scalability
 - Ability to add more service blades and redistribute traffic to new resource

Bucket-based LB'ing → Seamless Redistribution on Failover



Load Balancing: Common Methods (1 of 2)

- Direct Association (forwarding table)
 - New flow is consciously assigned to a specific target
 - New flows can be assigned to least-loaded server blade
 - Choice of target can be based upon current spare capacity and existing level of commitments
 - DPI packet blade is used to make service-related targeting decision
 - Software can dig deep into packet to find key attributes

Issues

- Compute-intensive – requires additional packet blades for LB'ing
 - Low density: roughly 1 packet blade per 10GbE link
- Introduces latency – packet stored & analyzed then forwarded
- Despite carefully targeting each new flow, server blades may still become unfairly overloaded due to future traffic patterns

Load Balancing: Common Methods (2 of 2)

■ Statistical

- A flow is predictably (but arbitrarily) always sent to a particular blade
- Choice of target typically based upon a hash of IP header
 - Fair distribution relies on statistical fairness across a large population set
- LB'ing function provided as an intrinsic feature of the switch
 - No additional components (system needs a switch anyway)
 - Algorithm implemented in switch silicon → low latency
 - Capable of supporting 8 x 10GbE links with 64Byte packets at wireline speeds*

*FlexCore ATCA-FM40 with Fulcrum FM4224 silicon

Issues

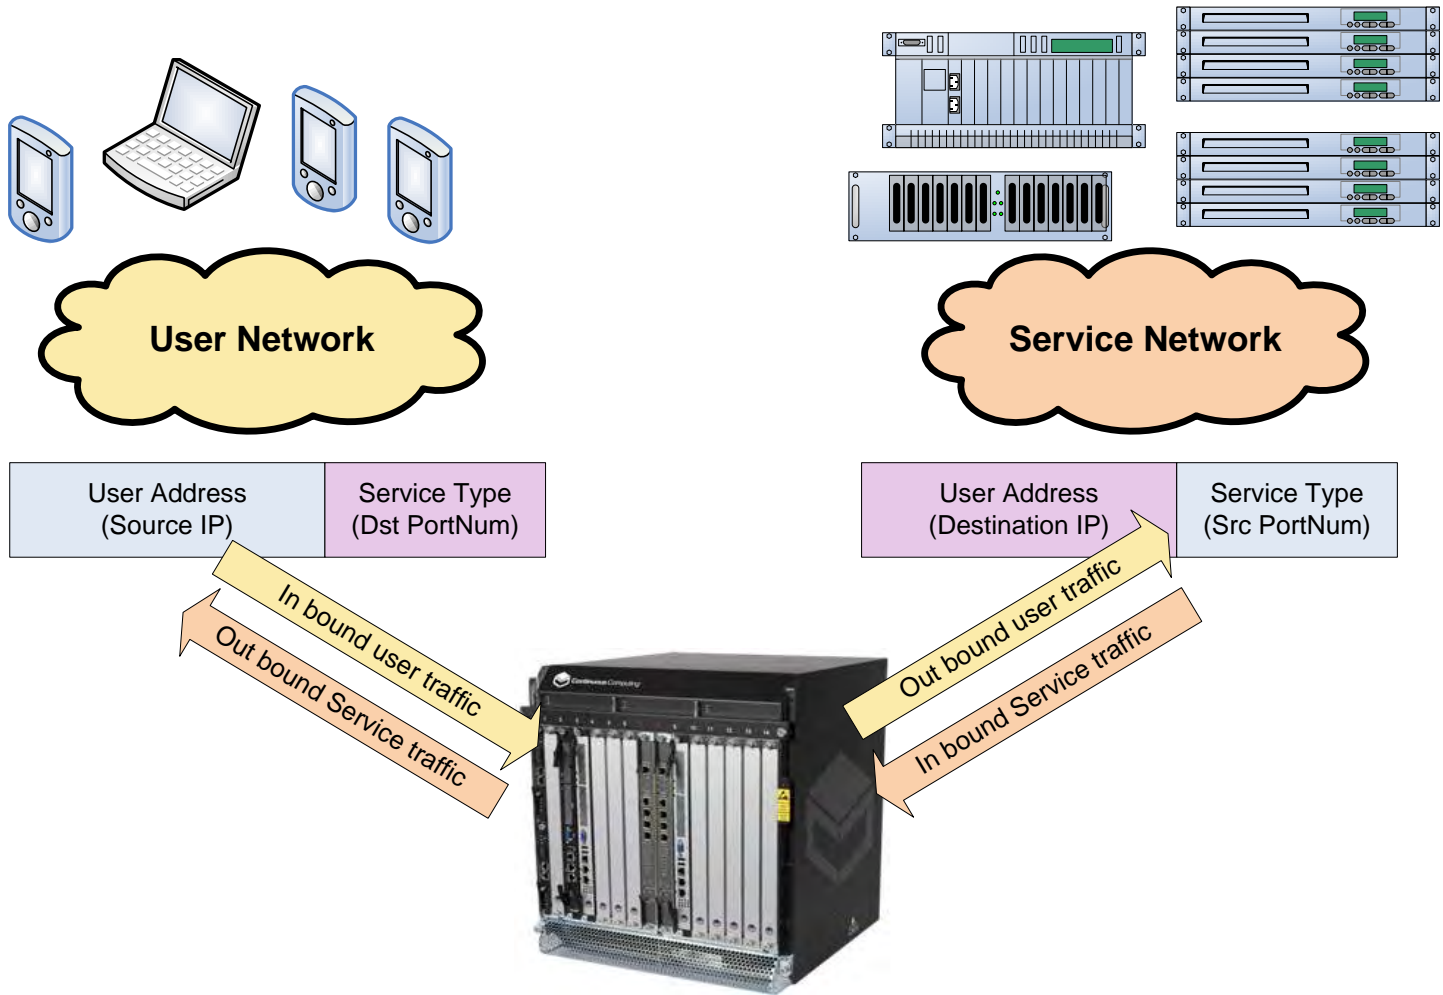
- Fair distribution relies upon statistical variance in flow attributes
- Difficult to locate 'user' IP header from deeply encapsulated packets

Load Balancing: Flows vs. Subscribers

- Most LB'ing solutions use 'flows' to determine target
 - 3-tuple: source IP, destination IP, protocol ID
 - 5-tuple: 3-tuple, source L4 port, destination L4 port
- Increases degree of statistical variance between flows
- Often done for convenience (re-use of LAG logic)

- However, most applications need to track a subscriber
 - All traffic relating to a particular subscriber (of all types and for all services) needs to be targeted through the same packet blade

Load Balancing by User IP Address



FlexCore ATCA-FM40: Load Balancing Advantages

- Leverages Fulcrum's FM4224 10GbE switch
- Frame Forwarding Unit (FFU) allows identification of user in both user- and network-originating packets
- FFU can also detect and drop (filter) packets of no interest
 - i.e., only forward http traffic to packet blades
 - Based upon L4 src/dst port number
- FFU can manipulate an added VLAN tag to fine tune the target to receive the packet
 - VID defines target → node_blade:processor:thread



Thank You

James.Radley@ccpu.com