



UCX Wire Compatibility

Mikhail Brinskii

UCF 2023

Agenda

- Introduction
- UCP protocols compatibility
- Endpoint reconfiguration issues
- CI pipeline improvements
- Work in progress

UCX Wire Compatibility

Introduction

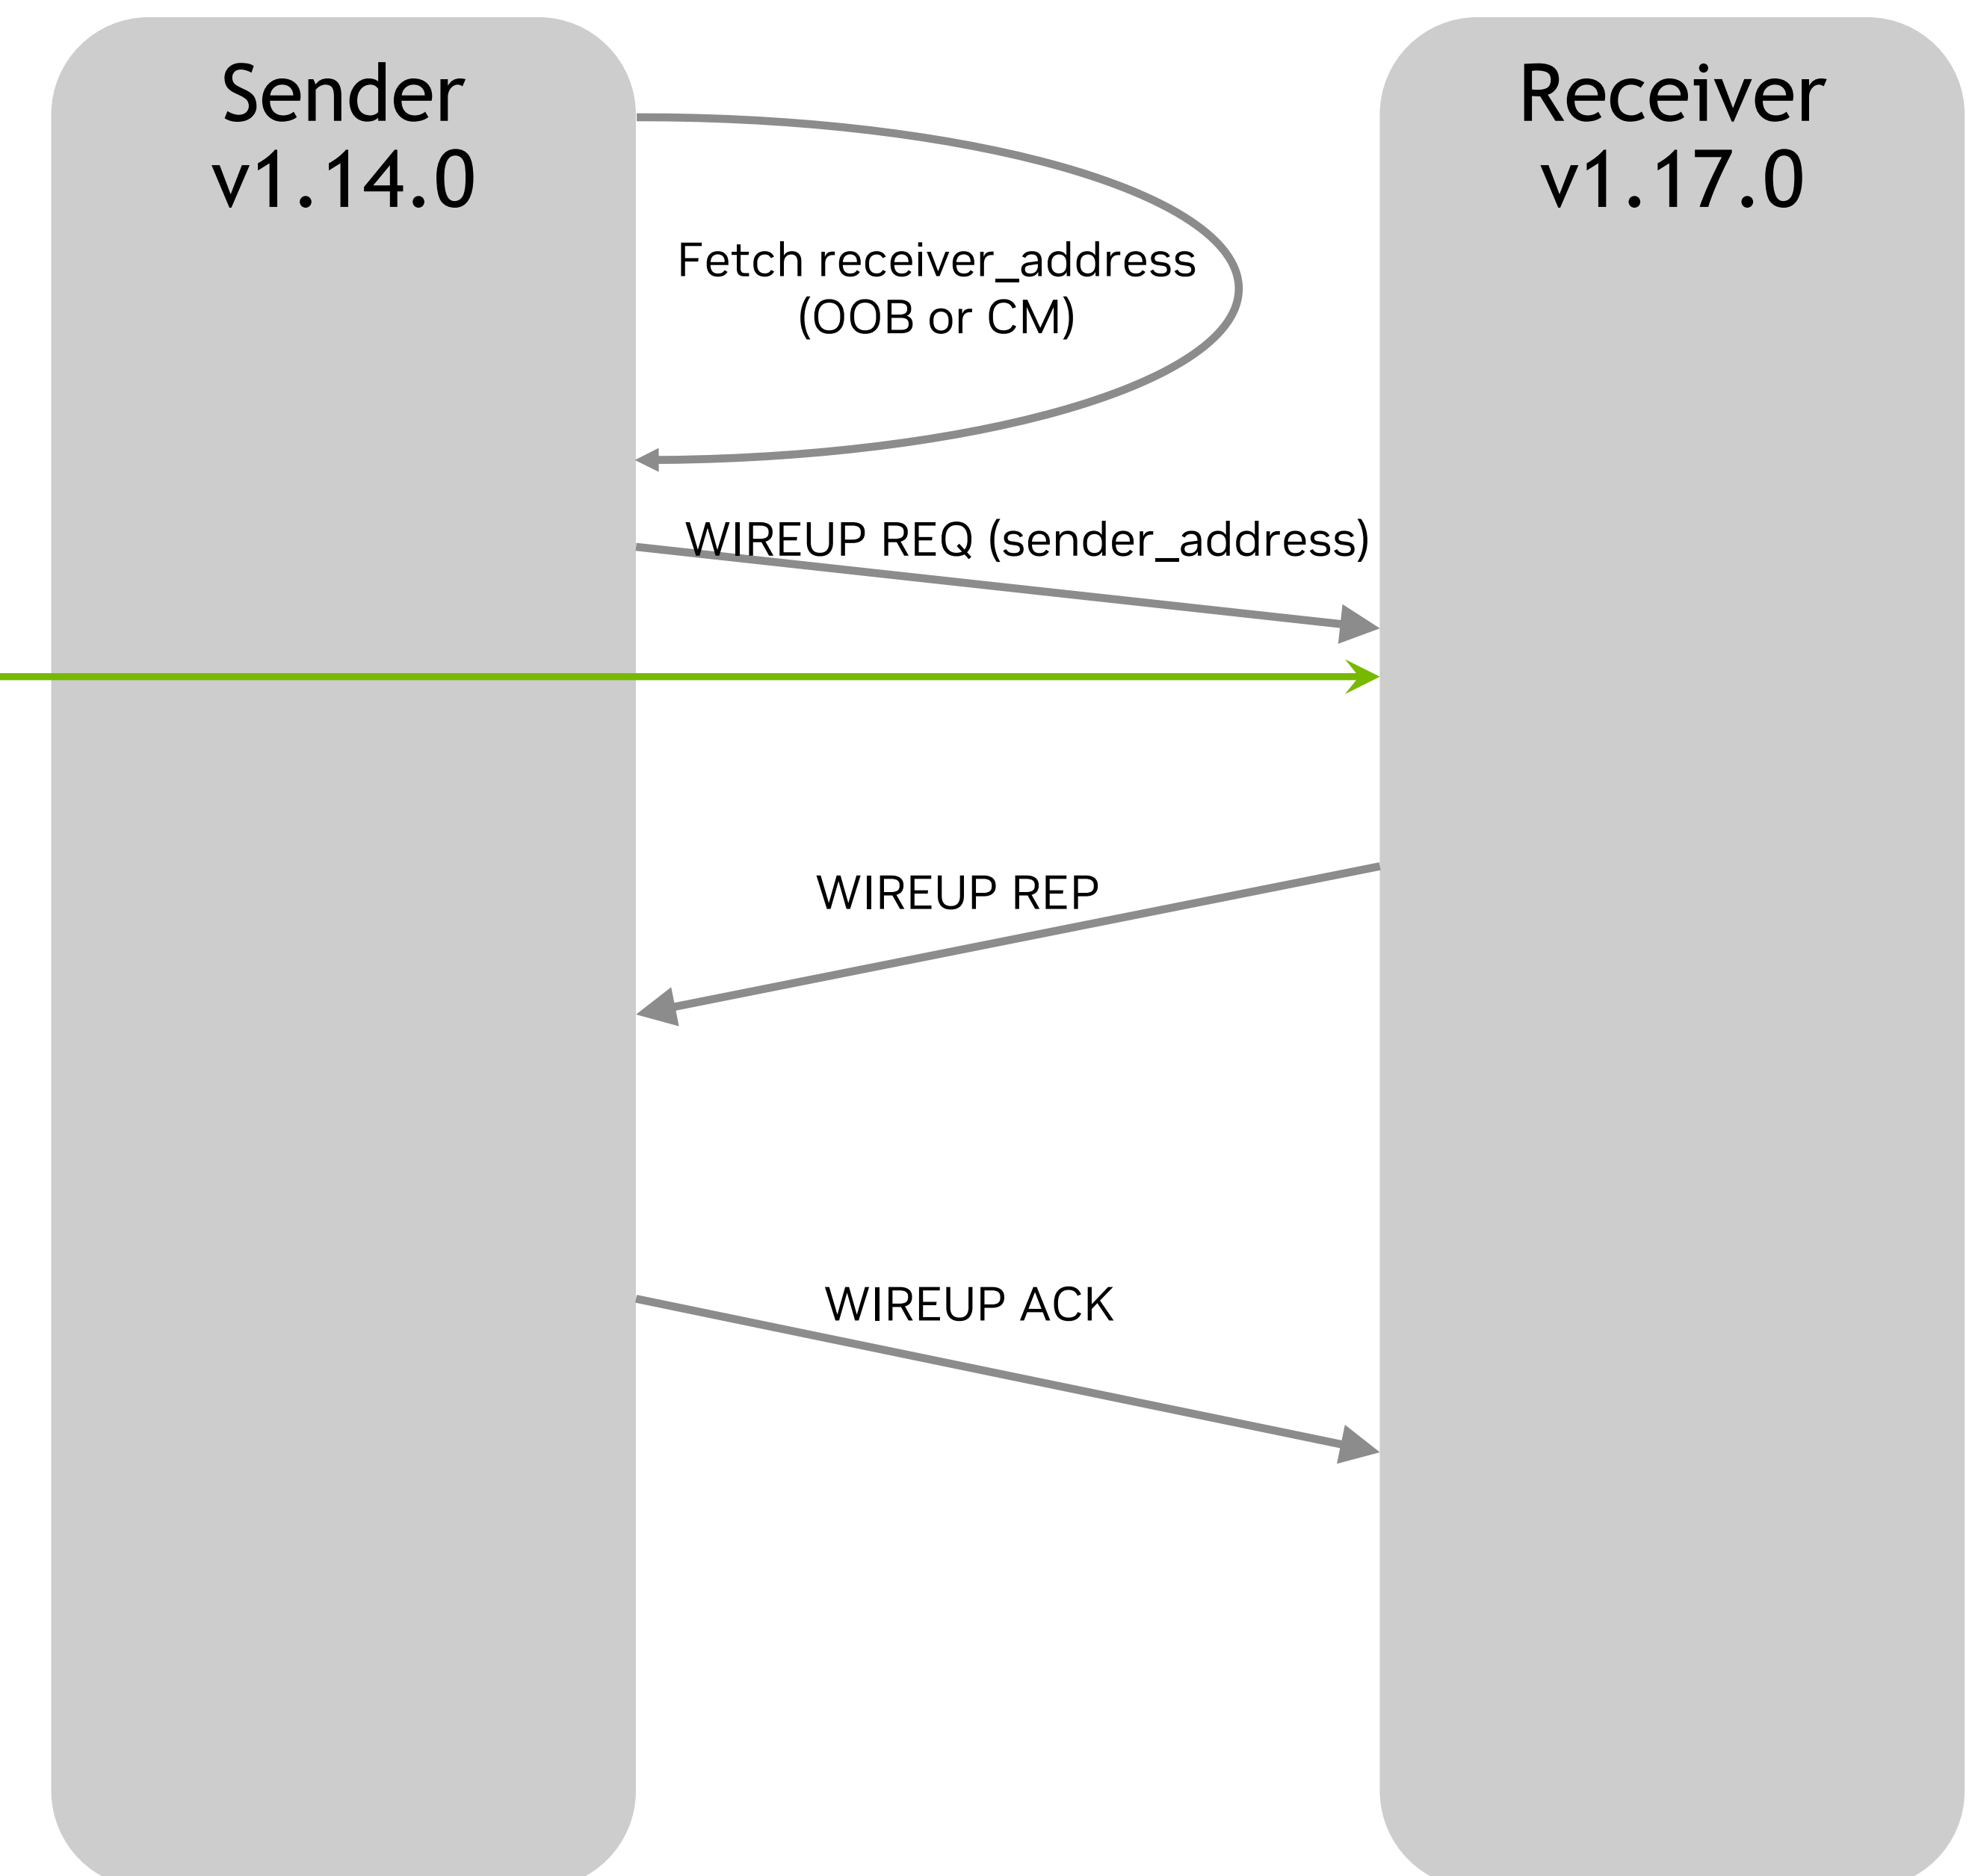
- Wire compatibility assumes that applications may use different versions of UCX library
- Important in cluster and cloud environments
- UCX is wire compatible since UCX v1.14.0
 - Limited compatibility since UCX v1.12.0 (only network transports and single device in CM flows)
- Currently up to four previous UCX versions supported
- Wire compatibility in UCX means:
 - Protocol headers are compatible
 - Connection establishment protocol is aligned
 - Configuration parameters have the same default values

UCX Wire Compatibility

UCP protocols compatibility

- Peers always exchange worker addresses
 - Via OOB mechanism for connect to worker address flows
 - Via CM lane for connect to IP address flows
- Worker address contains UCX minor version
 - Packed in 4 bits, assuming no more than 8 previous versions supported
 - Starting from UCX v1.16
 - Older versions are packed as version 0 (i. e. legacy)

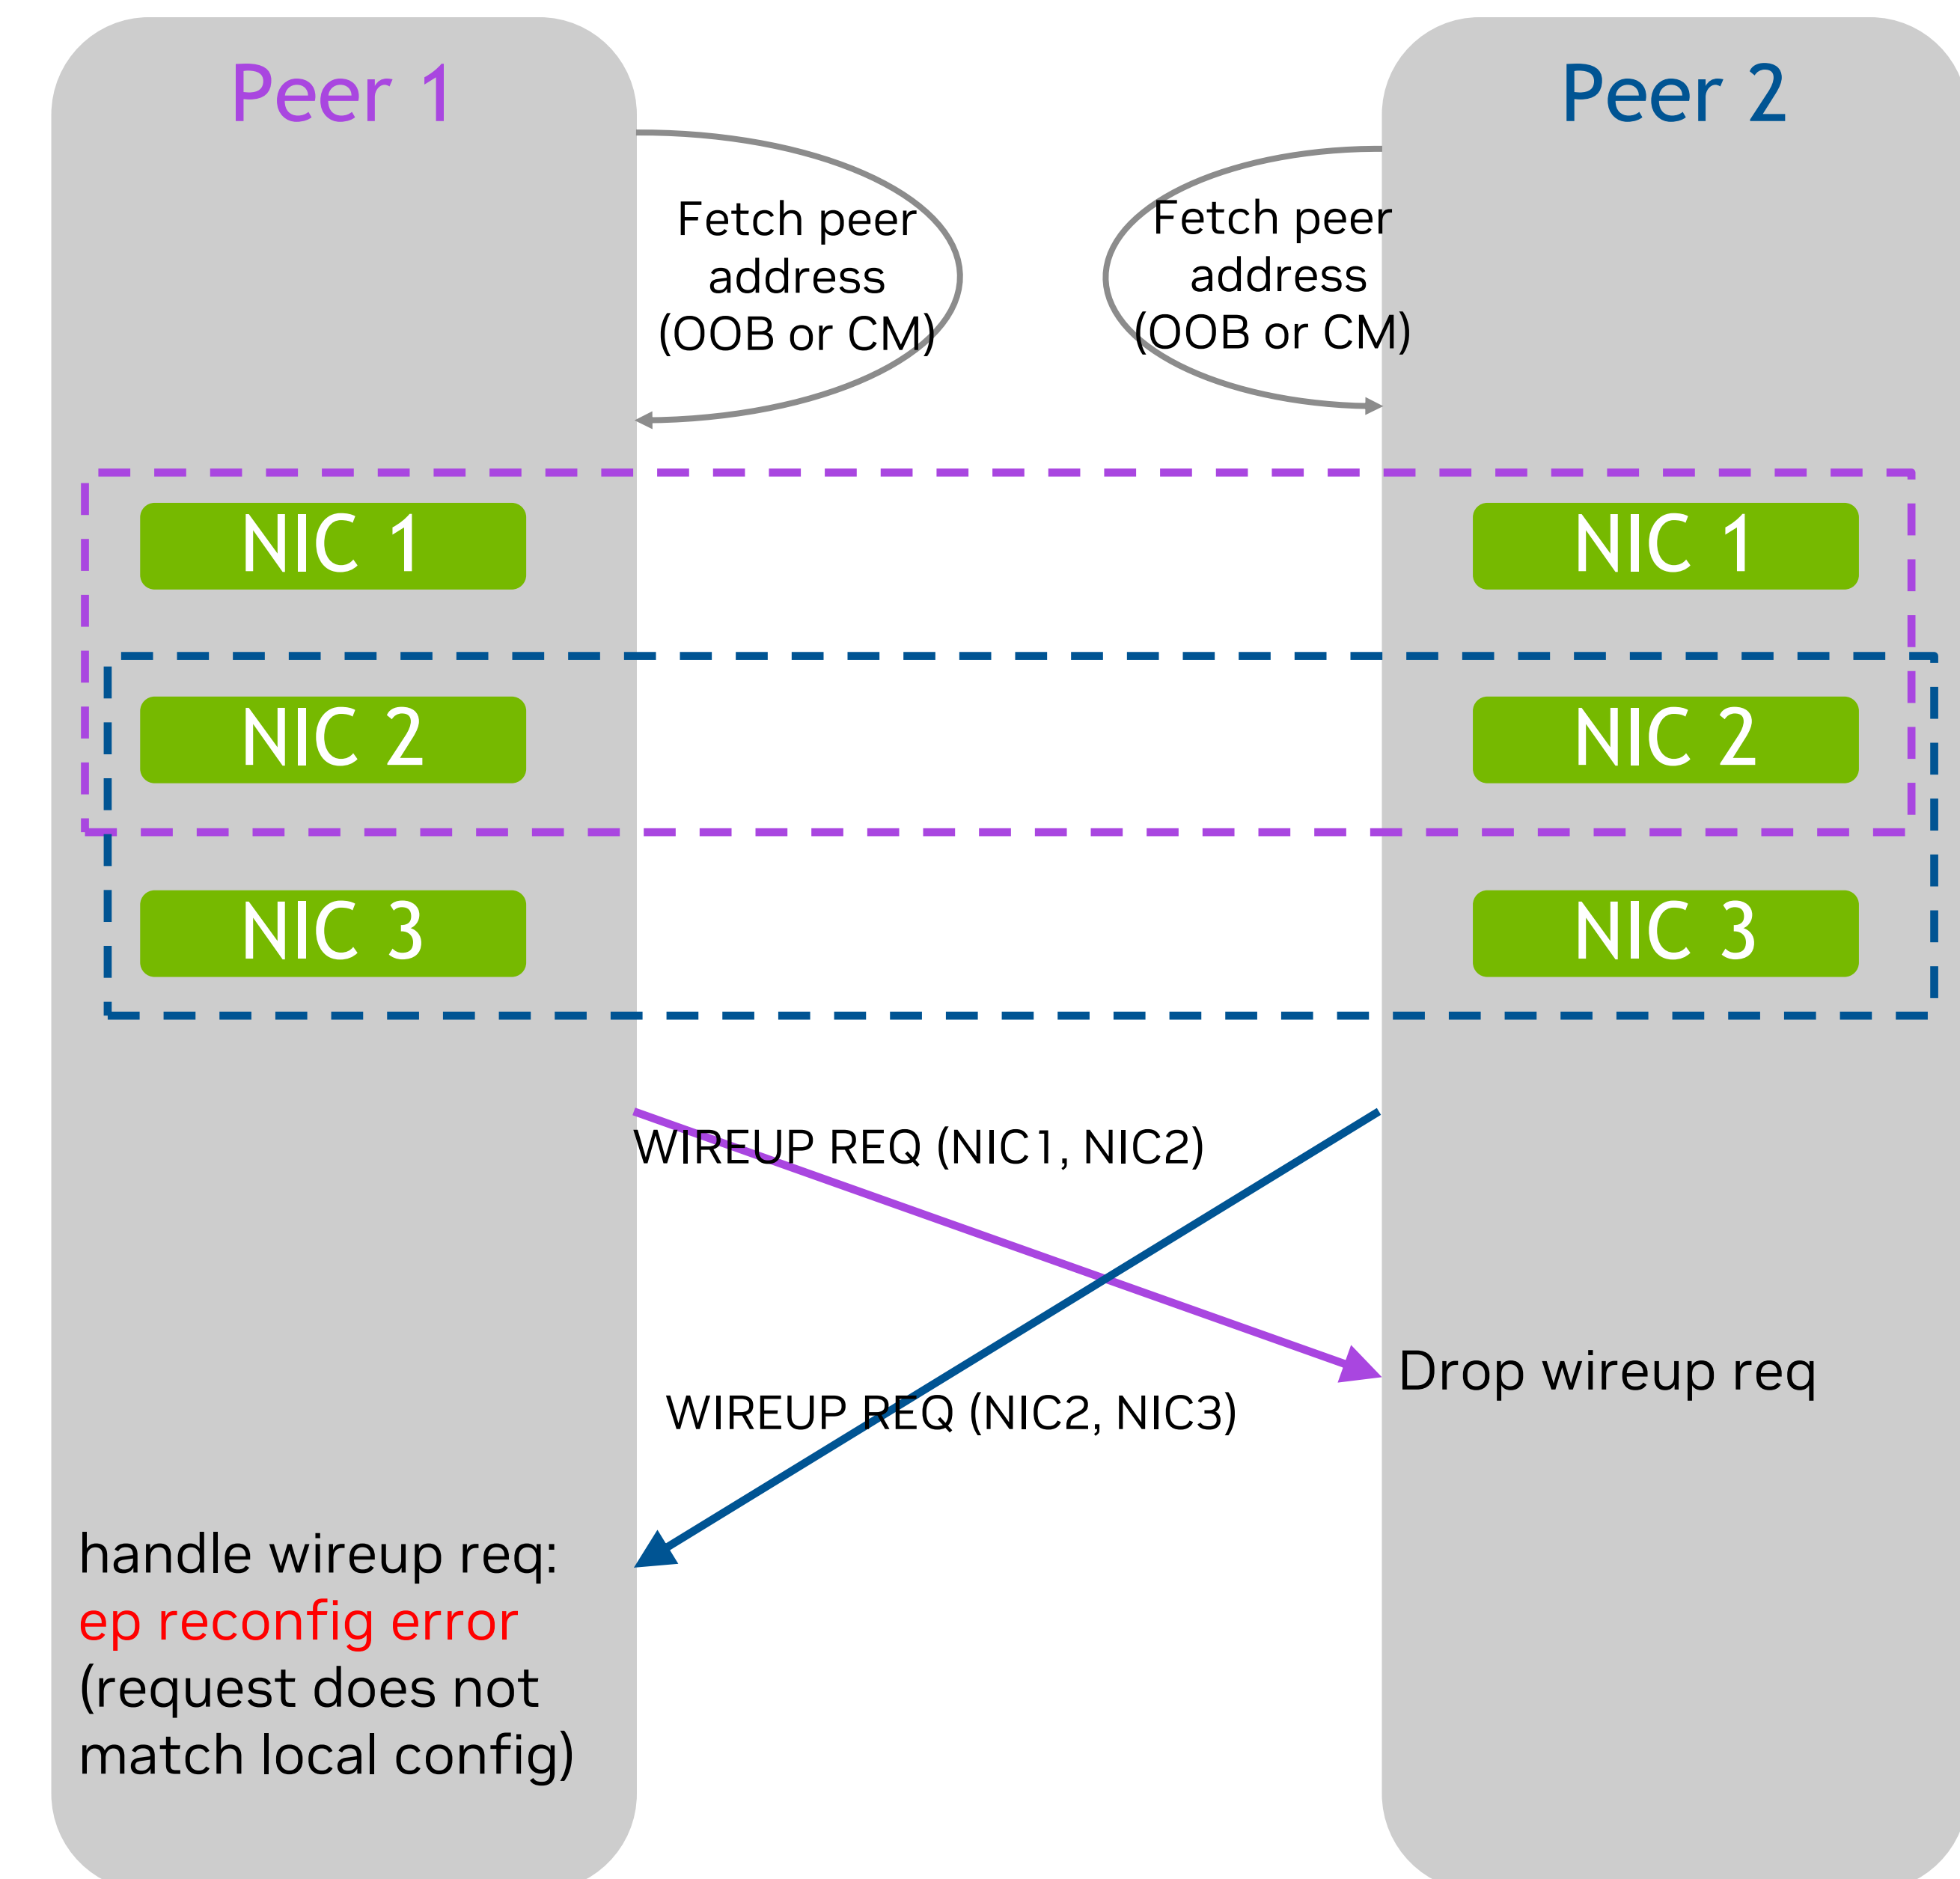
```
if (sender_address.version < 17) {  
    initialize_old_protocols();  
} else {  
    /* Peer can parse newest protocol  
    * headers!  
    */  
    initialize_new_protocols();  
}
```



UCX Wire Compatibility

Endpoint reconfiguration issues

- Peers may initiate connection establishment simultaneously
- Connection establishment:
 - Select communication lanes based on local performance characteristics and initial peer address
 - Send wireup request with address containing selected lanes information
- Potential issues in head-to-head resolution:
 - Peer1 creates initial endpoint configuration
 - Peer2 also creates initial endpoint configuration
 - These configurations are different due to some reason (e. g. some performance value is calculated differently)
 - Peer1 attempts to handle wireup request from Peer2 but fails because the address in this request does not match the local endpoint configuration.



UCX Wire Compatibility

Endpoint reconfiguration issues

- Endpoint reconfiguration symptoms:

```
wireup.c:1554 Fatal: endpoint reconfiguration not supported yet

wireup.c:1204 UCX ERROR old: am_lane 0 wireup_msg_lane 6 cm_lane <none> keepalive_lane <none> reachable_mds 0x3fe
wireup.c:1214 UCX ERROR old: lane[0]: 7:sysv/memory.0 md[2] -> md[2]/sysv/sysdev[255] am am_bw#0
wireup.c:1214 UCX ERROR old: lane[1]: 30:xpmem/memory.0 md[9] -> md[9]/xpmem/sysdev[255] rkey_ptr
wireup.c:1214 UCX ERROR old: lane[2]: 11:rc_mlx5/mlx5_0:1.0 md[4] -> md[4]/ib/sysdev[255] rma_bw#0
wireup.c:1214 UCX ERROR old: lane[3]: 26:rc_mlx5/mlx5_4:1.0 md[7] -> md[7]/ib/sysdev[255] rma_bw#1
wireup.c:1214 UCX ERROR old: lane[4]: 29:cma/memory.0 md[8] -> md[8]/cma/sysdev[255] rma_bw#2
wireup.c:1214 UCX ERROR old: lane[5]: 16:rc_mlx5/mlx5_2:1.0 md[5] -> md[5]/ib/sysdev[255] rma_bw#3
wireup.c:1214 UCX ERROR old: lane[6]: 21:rc_mlx5/mlx5_3:1.0 md[6] -> md[6]/ib/sysdev[255] rma_bw#4 wireup
wireup.c:1204 UCX ERROR new: am_lane 0 wireup_msg_lane 3 cm_lane <none> keepalive_lane <none> reachable_mds 0x3fe
wireup.c:1214 UCX ERROR new: lane[0]: 7:sysv/memory.0 md[2] -> md[2]/sysv/sysdev[255] am am_bw#0
wireup.c:1214 UCX ERROR new: lane[1]: 30:xpmem/memory.0 md[9] -> md[9]/xpmem/sysdev[255] rkey_ptr
wireup.c:1214 UCX ERROR new: lane[2]: 11:rc_mlx5/mlx5_0:1.0 md[4] -> md[4]/ib/sysdev[255] rma_bw#0
wireup.c:1214 UCX ERROR new: lane[3]: 26:rc_mlx5/mlx5_4:1.0 md[7] -> md[7]/ib/sysdev[255] rma_bw#1 wireup
```

- Endpoint reconfiguration GH issues:

- <https://github.com/openucx/ucx/issues/9159>
- <https://github.com/openucx/ucx/issues/4160>
- etc

UCX Wire Compatibility

CI pipeline improvements

- New stage in CI PR pipeline
- Run `ucp_client_server` and `ucp_hello_world` examples:
 - With different UCX versions
 - On different setups
 - With different transports (network or shared memory)
- 1-line change is needed to add testing of new UCX version
- Wire compatibility stage is defined in:
 - <https://github.com/openucx/ucx/blob/master/buildlib/pr/main.yml>
 - https://github.com/openucx/ucx/blob/master/buildlib/pr/wire_compat.yml
 - https://github.com/openucx/ucx/blob/master/buildlib/tools/test_wire_compat.sh

The screenshot displays a CI pipeline interface. On the left is a sidebar with navigation options: Boards, Repos, Pipelines, Pipelines, Releases, and Artifacts. The main area shows a list of jobs under the 'WireCompat' stage. A green box highlights the 'WireCompat' stage name. The jobs listed are:

Job Name	Duration
Build on althca	5m 18s
Build on gpu	4m 0s
Build on new	4m 24s
Build on bond	4m 49s
Test wire compatibility althca...	3m 51s
Initialize job	7s
Checkout openucx/ucx@refs/...	13s
Download UCX current	9s
Build ucx artifact	2m 37s
ucp_client_server with UCX v1...	10s
ucp_hello_world with UCX v1...	10s
ucp_client_server (no sm) with...	10s
ucp_hello_world (no sm) with ...	10s
Post-job: Checkout openucx/uc...	1s
Finalize Job	<1s
Test wire compatibility althca...	3m 48s
Test wire compatibility althca ...	4m 4s

On the right, a terminal window shows the output of the test jobs. The output includes:

```
26 [1701338041.701813] [sw
27 [1701338041.710844] [sw
28 [1701338041.710868] [sw
29 [1701338041.710949] [sw
30 [1701338041.715738] [sw
31 [1701338041.715752] [sw
32 Waiting for connection.
33 server is listening on
34 + LD_LIBRARY_PATH=/scrap
35 + stdbuf -oL /scrap/azu
36 [1701338046.742496] [sw
37 [1701338046.751088] [sw
38 [1701338046.751100] [sw
39 [1701338046.752270] [sw
40 [1701338046.752270] [sw
41 [1701338046.752279] [sw
42 [1701338046.755991] [sw
43 Server received a conne
44 [1701338046.760345] [sw
45 [1701338046.768898] [sw
46 [1701338046.778169] [sw
47 [1701338046.779013] [sw
48 Client: iteration #1
49
50
51 -----
52
53 ABCDEFGHIJKLMNO.
54
55
56 -----
57
58 Server: iteration #1
59 UCX data message was re
60
61
```

UCX Wire Compatibility

Work In Progress

- Enhance wire compatibility CI testing: [UCX PR #9460](#)
 - CI check that different UCX versions use the same performance values for lanes selection
- Support endpoint reconfiguration: [UCX PR #9234](#)
 - Full support of head-to-head resolution of wireup requests
 - Update local endpoint configuration to match address in incoming wireup request

