

# UCX Shared Memory in Containers



# UCX Shared Memory (current state)

- Intra-node host memory transports:
  - SYSV
  - POSIX
  - CMA
  - KNEM
  - XPMEM
- All transports construct device address based on MAC address
- Check for peer reachability using device address fails if peers use different NET namespaces

# UCX Shared Memory (namespace requirements)

Transport   NS	IPC	MNT	NET	PID	USER	UTS
POSIX	X			X		
SYSV	X					
CMA	X			X		
KNEM	X					
XPMEM	X					

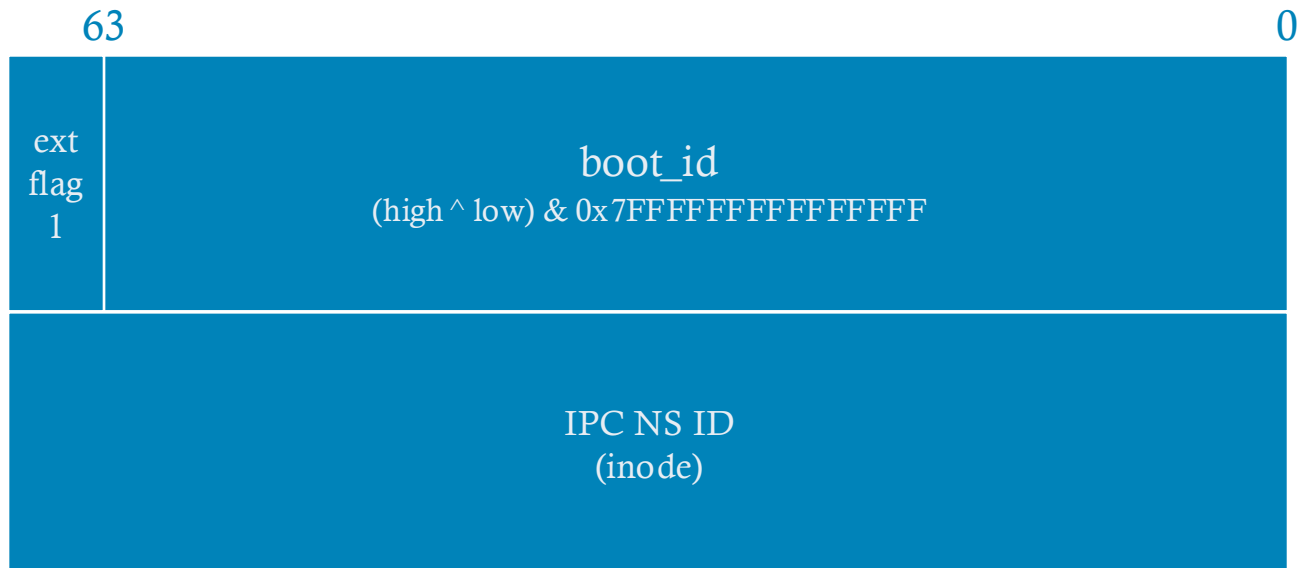
- Need to distribute IPC and PID namespace IDs for correct reachability check
- But need to avoid address growth where it is possible

# UCX Shared Memory Device Address

- Default (root IPC NS):



- Extended (non root IPC NS):

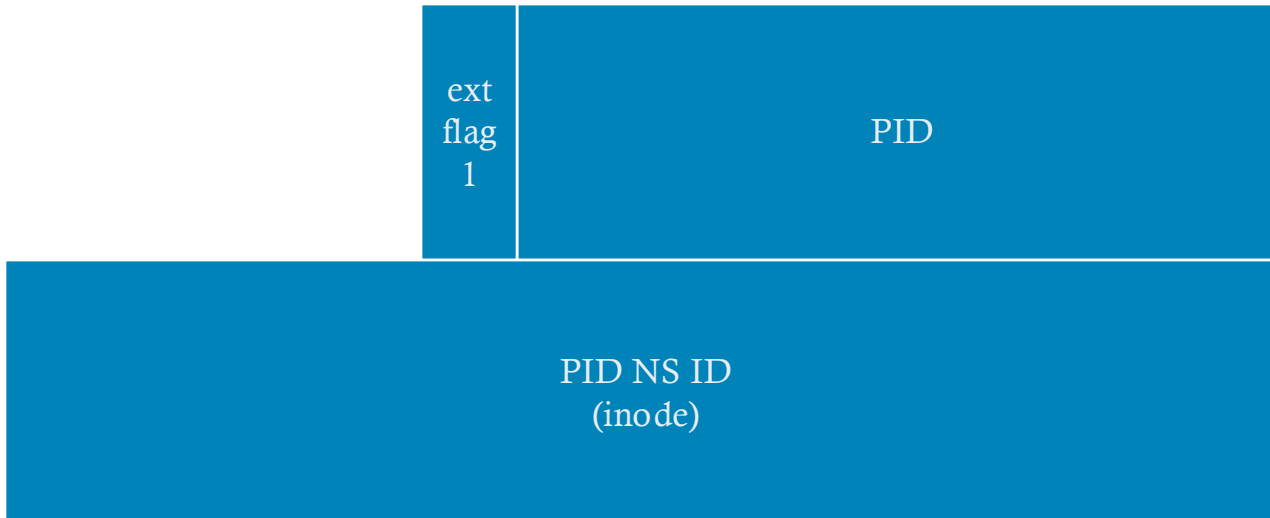


# UCX CMA Interface Address

- Default (root PID NS):



- Extended (non root PID NS):



# UCX POSIX Interface Address

- PID namespace id is added to POSIX interface address, if:
  - procs is used for sharing POSIX backing file (using path /proc/PID/fd/FD)
  - PID namespace is not root
- In this case interface address is:

