# Extending ZeroMQ and/or NanoMSG with support for SCIF

› **26/02/2016**

**ALICE CWG13 Meeting**


ICE-DIP is a European Industrial Doctorate project funded by the European Community's 7th Frameworkprogramme Marie Curie Actions under grant PITN-GA-2012-316596

# The high level idea in a nutshell

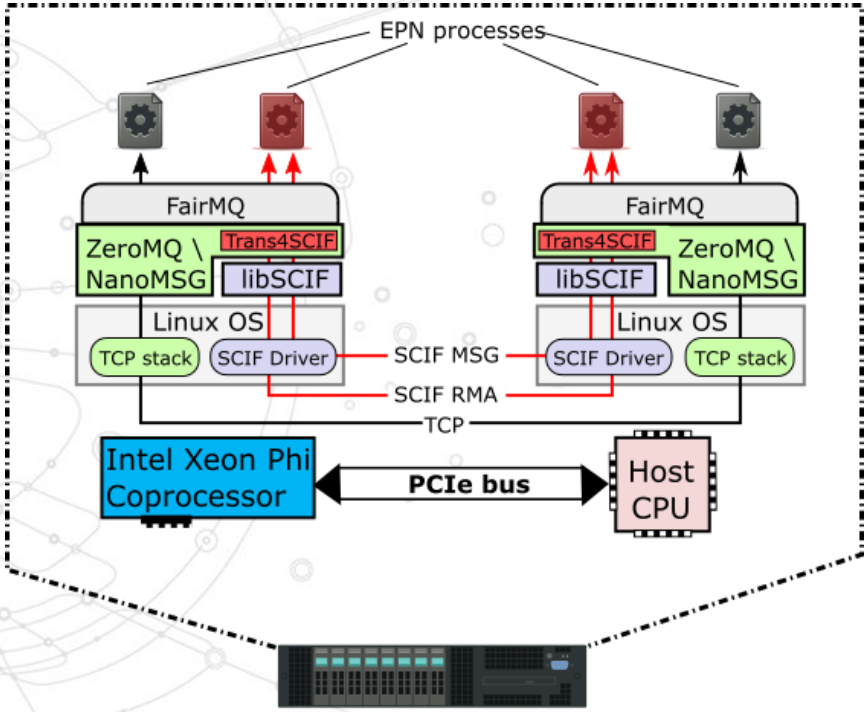


- Complete EPN processes on Xeon Phi
- Transparently (No API change)
- Efficiently (SCIF)



Background image: Shutterstock

# Three possible (?) solutions

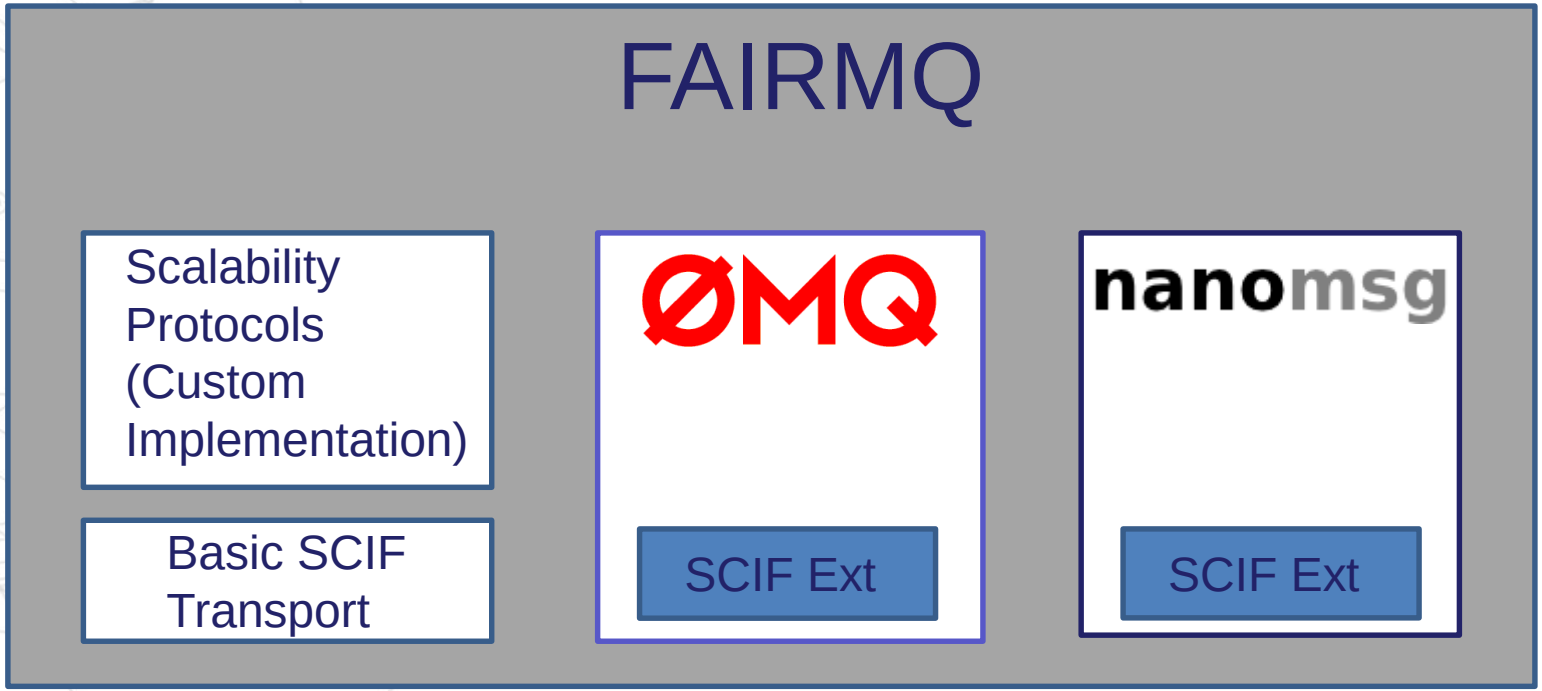




Background image: Shutterstock

# High cost (Development and Maintenance)

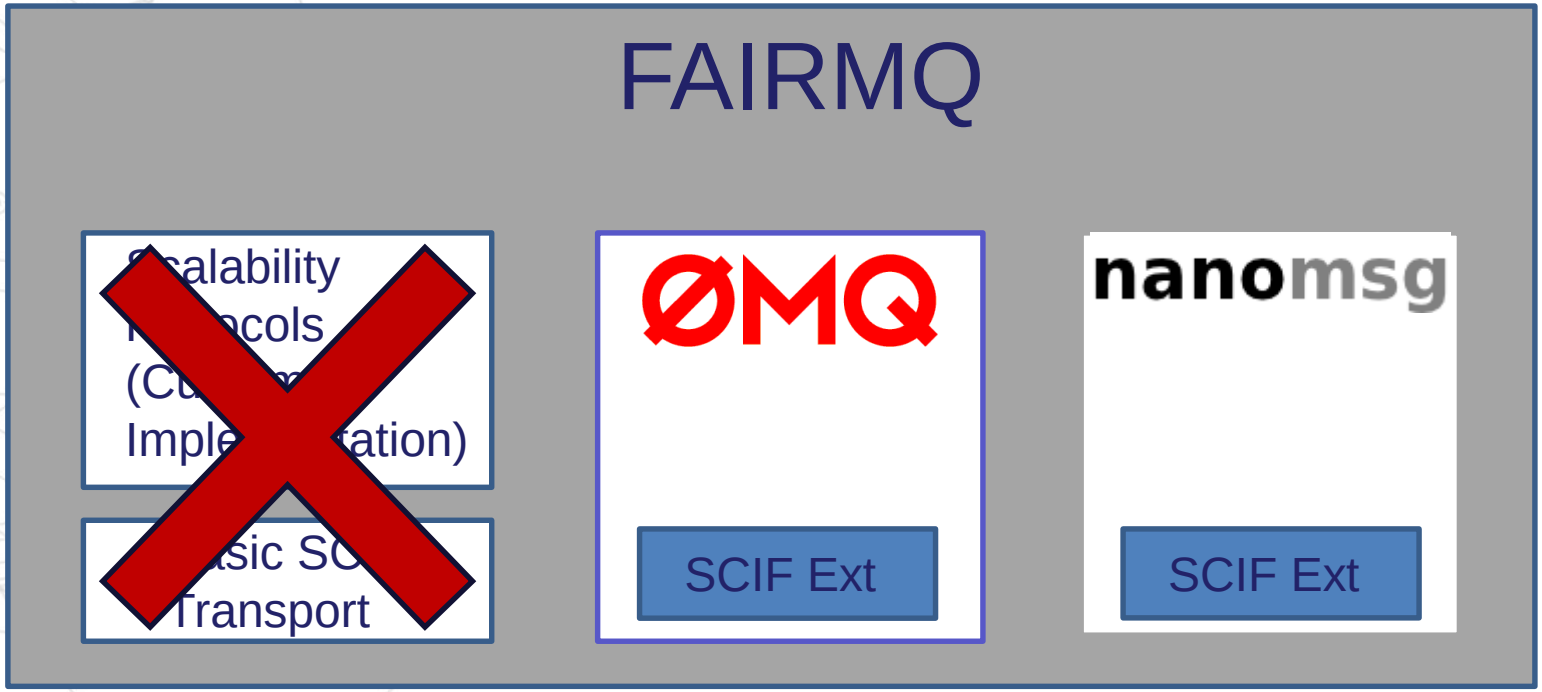




Background image: Shutterstock

# Bad investment (also I've tried and failed)

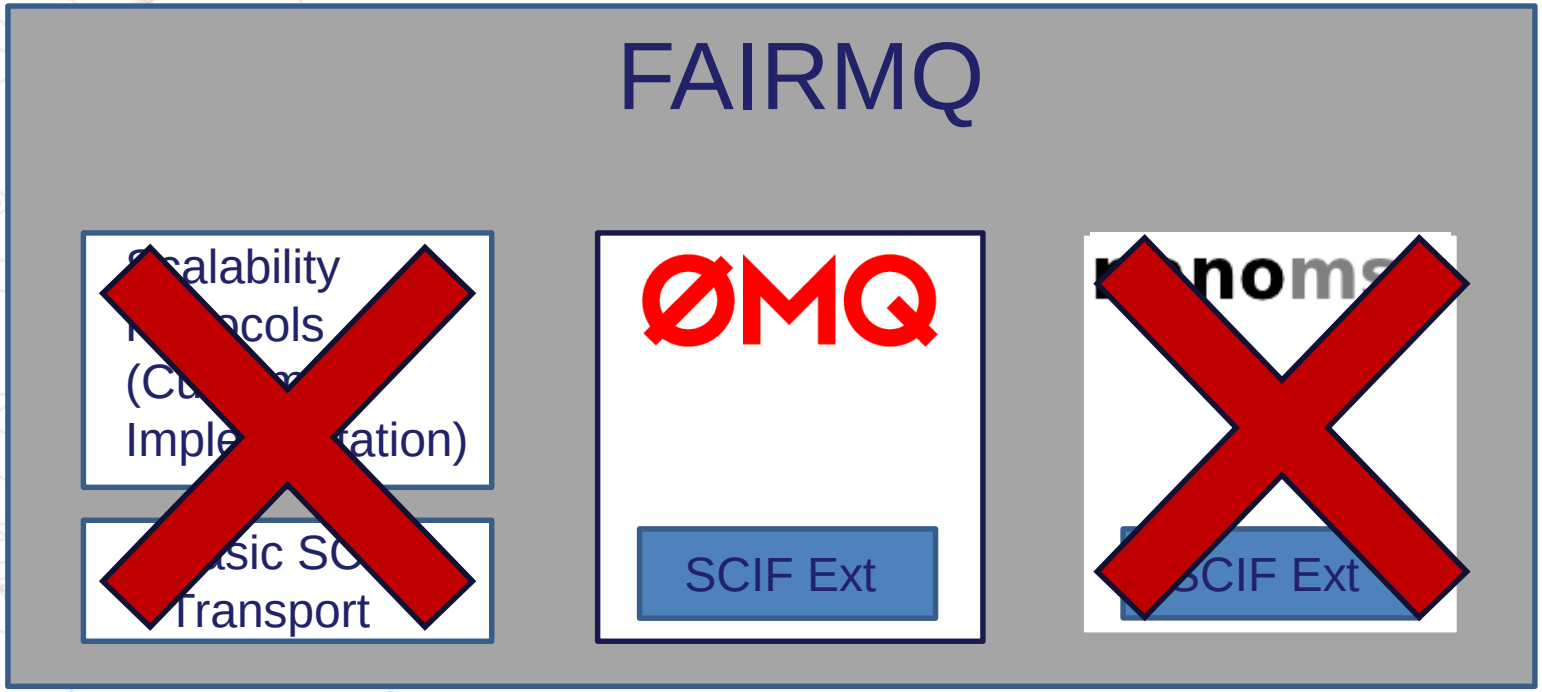




Background image: Shutterstock

# ZeroMQ extension de-mystified

- Nobody was able to add new transports



Background image: Shutterstock

# ZeroMQ extension de-mystified

- ~~Nobody was able to add new transports~~
  - VMCI support was added 2 months ago
  - NORM support a year ago

# ZeroMQ extension de-mystified

- ~~**Nobody was able to add new transports**~~
  - VMCI support was added 2 months ago
  - NORM support a year ago
- **Complex codebase (practically impossible to extend it)**

# ZeroMQ extension de-mystified

- ~~**Nobody was able to add new transports**~~
  - VMCI support was added 2 months ago
  - NORM support a year ago
- ~~**Complex codebase (practically impossible to extend it)**~~
  - I implemented a basic SCIF extension in a **week!**

# The status of the extension

- To enable it `./configure --with-scif`
- You can run perf tests

```
./local_thr scif://6666 1000 1000
./remote_thr scif://0:6666 1000 1000)
```

- There is a BUG: the connector socket can't receive.
- The scif_engine doesn't use RDMA (yet!)

# How to extend ZeroMQ with transports

- **Implement 4 modules**
  - Address (myaddr:// address resolution)
  - Connecter
  - Listener
  - Engine (send/recv)
- **~100 lines of boilerplate glue code**

# Why not NanoMSG ?

- Was never used in production (to my knowledge)
- Development is stalled (community issues)
- Adding features on unstable code base (looking for trouble)
- My initial attempt failed (undefined behaviour of my implementation)

# Future perspective

- Wrap-up the SCIF extension by end of March
- Write documentation-publication in by mid-April and pull request for upstream
- FairMQ on Xeon Phi porting/tests
- Prepare for “Knights Landing” (Omni-path support in ZMQ etc.)

# Thank you


ICE-DIP is a European Industrial Doctorate project funded by the European Community's 7th Frameworkprogramme Marie Curie Actions under grant PITN-GA-2012-316596

CERN openlab

Background image: Shutterstock