

# TensorFlow Networking Status

2018.10.31

## Context Constraints

Need to support a wide range of OSs and platforms, only gRPC assumed.

- Clouds with preemption
- Dedicated HPC platforms
- Informal clusters
- little common low-level infrastructure

## General Characteristics

### Extension by C++ Inheritance and Registration

- GrpcServer: exports MasterInterface, WorkerInterface, can customize various components
- BaseRemoteRendezvous, BaseRecvTensorCall: facilitate remote Tensor retrieval
- WorkerInterface: RPC on protocol buffers. Any method can be overridden, e.g. GrpcRecvTensorAsync
- CollectiveExecutorMgr (new)

### Hooks for customization in GrpcServer

- Registration of servers by protocol
- RendezvousMgr creator function
- Worker creator function
- Service function
- CollectiveExecutorMgr function

Little or none of this is an official API, all subject to change at any time. In practice there are only a few extensions, has been fairly stable.

The only performance critical parts are RecvTensor (or RemoteRendezvous methods), and going forward RecvBuf (or CollectiveExecutorMgr related methods).

## Possible Issues to Address

1. Desire to eliminate /contrib in favor of an official API with 3rd party plugins. Is that something other than freezing and blessing some part of the currently used virtual interfaces?
  - a. Simplify end-user configuration experience
  - b. Presubmit testing of changes is hard
2. Reduce the need to recompile everything with bazel. Would this be worth the trouble?
3. Some kind of (semi-)official MPI or ib/verbs support. What would that look like?
4. How does NCCL or a similar non-NVIDIA interconnect related utility fit in?