



Pilot: IPv6 only Grid

Summary

- Goals of pilot
- Limitations / hacks / corners cut
- Integration with internal infrastructure
- HTCondor routes & config
- How experiments can participate

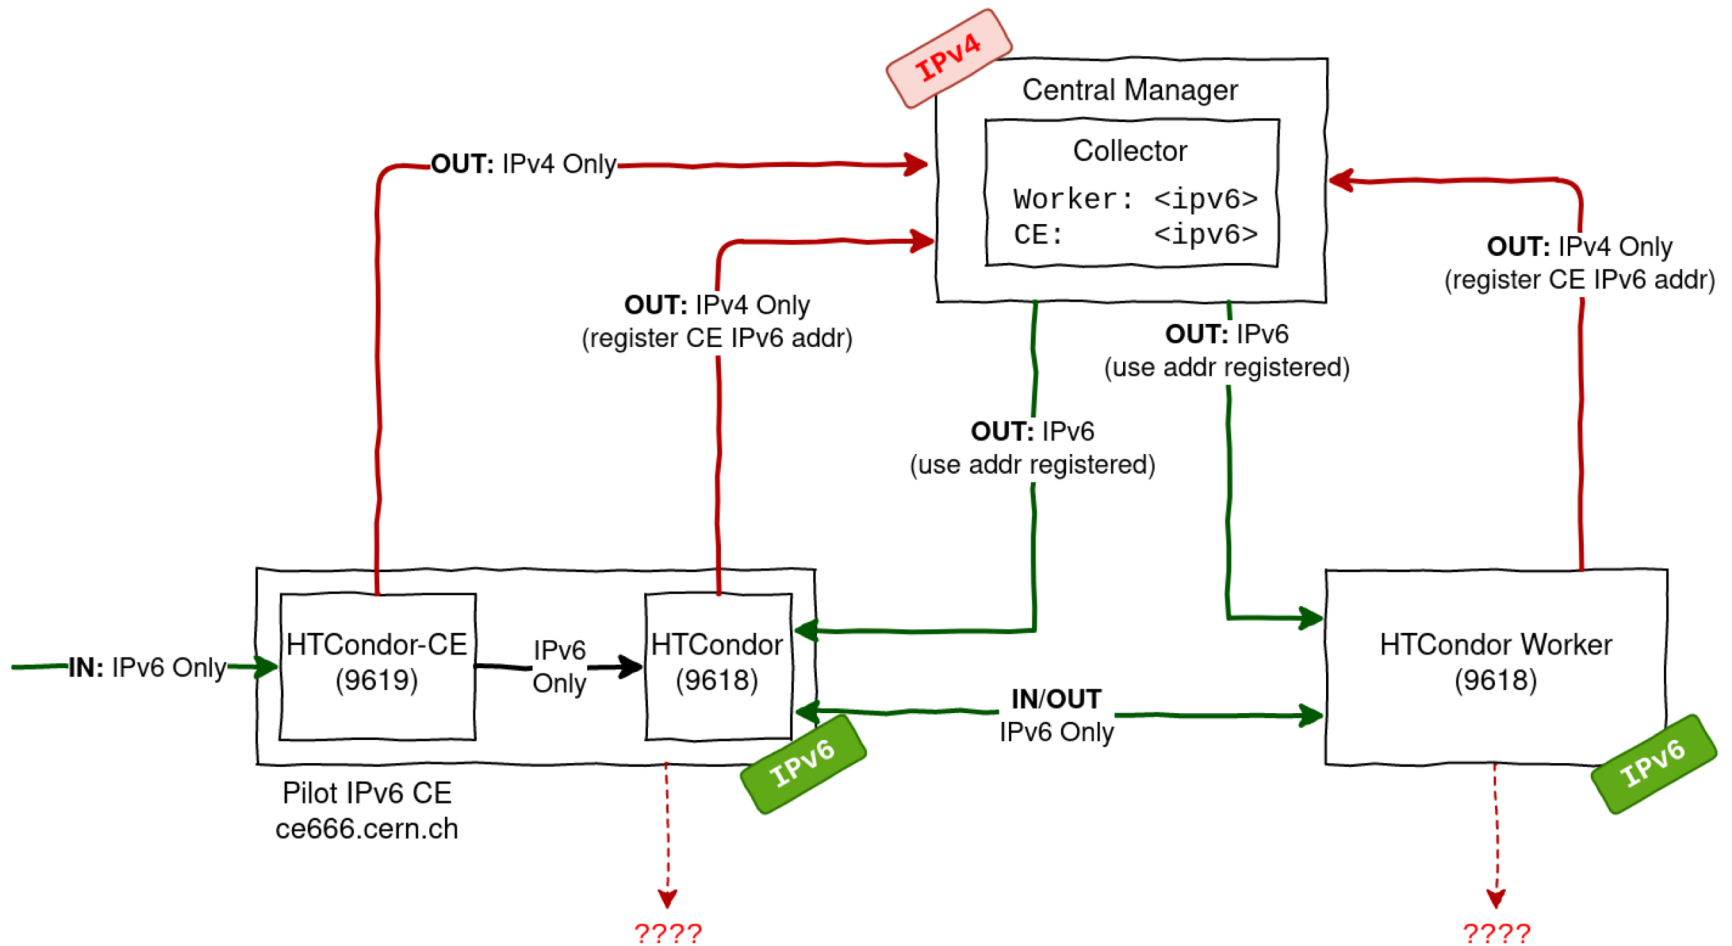
Goals of Pilot

- Can we find current issues with IPv6?
 - We can wait till everything is fixed, or we can just see what's broken
- Gain experience running IPv6 in production
 - Even CERN has some IPv4 pressure
 - Can next DC be IPv6 only
- End-to-end IPv6: submission to execution
 - IPv6 only workers primary goal, but investigating CEs too
- **Not necessary to wait for perfect IPv6 infra to make progress**

Limitations / Hacks

- Available networks mean both the workers and CEs are dual stack
 - iptables rules to block & log IPv4
 - ensure we are configured to prefer IPv6
- Various system services require IPv4
 - where we can't configure to use IPv6 we are allowing specific IPv4 traffic
 - one goal is to discover these limitations, but the point is to be able to run jobs

HTCondor traffic



Integration with services

Service	Protocol	Issue
Argus	IPv4 Only	Needs configuration
MONIT	IPv4 Only	MONIT alarm infrastructure currently IPv4 (working on dual stack)
OpenStack metadata	IPv4 Only	Support added for IPv6 in Nova (Victoria). Tools like facter will need to be updated
SSSD	IPv4/IPv6	Needs the configuration flag <code>lookup_family_order=ipv6_first</code> to pickup IPv6
Chrony	IPv4/IPv6	Works, but not all time servers configured dual stack
BDII	IPv4	BDII queries HTCondor-CE port 9619 via IPv4
fetch-crl	IPv4/IPv6	Needs “inet6glue” config param and dependency
CVMFS	IPv4/IPv6	Needs <code>CVMFS_IPFAMILY_PREFER=6</code>

HTCondor Configuration

- We DENY IPV4 traffic to 9619 (CE) and 9618
- PREFER_IPV4=False
 - Needed on CE & on workers
- ENABLE_IPV4=False on CE
 - Even with no IPV4 in DNS, the CE MyAddress will have IPV4 address without this setting
 - `_condor_PREFER_IPV4=False` env var for clients can workaround, otherwise global settings on schedds etc

HTCondor-CE routes

- CE defines an IPv6 only route with:

```
set_Requirements = (IPV6Only =?= True);  
    set_WantIPv6Only = True;  
    Requirements = (TARGET.WantIPV6Only =?= True ||  
TARGET.queue =?= "IPv6");
```

- Worker node requires jobs with an `WantIPv6Only` attribute with a `START` expression:

```
START = $(START) && (WantIPv6Only =?= True)
```

- Both sides of negotiation have requirement defined

Participate!

- CE: `ce666.cern.ch`
 - Only listening to IPv6
- JDL must have either:
 - `+WantIPv6Only=True`
 - `remote_queue="IPv6"`
- 158 cores for now
- We're logging IPv4 from the worker