CA Gateway Update

Ralph Lange
BESSY II
EPICS Meeting at KEK, May 2000

1 Overview *What is it? How does it work?*

The CA Gateway is an EPICS extension that was developed at the APS by Jim Kowalkowski. Further development was done by Janet Anderson (APS), Ken Evans (APS), Jeff Hill (LANL) and Ralph Lange (BESSY).

The Gateway is a proxy server for the Channel Access Protocol. It is an application that uses a CA client on the lower interface to connect to channels (PVs) and sets up monitors to those channels. It uses the portable Channel Access Server (CAS) to provide these channels to other CA clients (the upper interface). The channels' data is cached. There is an EPICS Access Security layer within the Gateway to configure whose clients from which hosts may have read or write access to a channel provided by the Gateway.

There are five basic functional blocks:

1. CA Server:
   The Gateway uses the portable Channel Access Server to provide its channels to the clients on the "upper" network.
2. Access Security:
   There is a level of configurable EPICS Access Security for the Virtual Connections (i.e. the channels the Gateway provides to its clients).
3. Virtual Connections:
   A Virtual Connection (VC) may either have a "real" PV name or an alias name that is mapped to an existing PV on the lower network. Configuration of Access Security, PV names and aliases is done using ASCII configuration files that are read by the gateway when it starts up.
4. Data Cache:
   For each requested "real" PV a connection via the lower network will be established. All channel attributes are read and a monitor for the value is set up. The attribute and value data is cached.
5. Channel Access Client:
   The CA client establishes the PV connections across the lower network. It is configurable through the usual set of environment variables

Key design features:

- There is only one connection between the Gateway and each server of “real” PVs (instead of one for each client). This minimizes the number of TCP connections as well on the upper as on the lower network.
- Read accesses from clients are answered from the data cache. There will be no traffic on the lower network for read accesses.
- The IOC sends monitor events (value or alarm severity changes) only to the Gateway, which distributes them to the clients. This minimizes the traffic on the lower network and the load on the IOC.
- PV connections are held open for a configurable time after the last client disconnects. This saves CA open or close actions on the IOC side if there are clients (like scripts or some SDDStools) that open, read and close a set of connections regularily.

2 Applications *What is it good for?*

Controlled Access to Secure Nets

Running the Gateway on a host equipped with two or more network interface cards may be used to isolate a private high reliability network from other networks. The network and IOC load implied from CA clients residing in other nets is limited and independent from the number and the behaviour of those clients. Access can be restricted using the Gateway’s Access Security Layer.

Channel Name Aliasing

A Gateway running on a host within the controls network may be used to introduce PV name aliases. This may be useful for testing database structures or to limit the IOC load for popular channels (like beam current, lifetime).

Panel Speed Up

If the OPI consoles are running on a fast network while the IOCs still use slow ethernet, accessing the PVs through a Gateway may speed up the connections to the OPI panels: The Gateway will handle the distribution of monitor events to many CA clients through much faster than the IOC, which has to send all the monitors using the slow line.

Local PV Repeater

The Gateway may be configured to use the loopback driver as upper interface. So all CA clients on this host may connect to the local Gateway, which acts as a repeater and bundles all CA connections on this host. (Like one Display Manager process bundles all CA connections from different panels on one display.)

3 Status *Latest improvements and bugfixes*

There have been some interesting improvements to the Gateway during the last months:

- Alarm Handler connections are handled properly. As the Alarm Handler is available on the WIN32 platform now, physicists and hardware people may configure and run Alarm Handlers on their desktops without putting additional load on the IOCs.
- The Gateway may be configured to completely deny access from certain hosts. This allows suppression of loops over multiple Gateways.

- The PV name patterns used in the configuration have been switched from simple filename glob patterns to regular expressions (GNU regex library). Alias commands may include sub-expressions in the “real name” part to allow very flexible alias name definitions.
- IOC reconnect situation are handled correctly. (Beacons are sent to the Gateway clients.)

The Gateway still needs a patched version of base, which complicates building the Gateway a lot. There is a drop-in tar that may be unpacked into a standard 3.13.2 base to make things easier. The next release of the Gateway will probably compile against a standard off-the-shelf version of base.

Release 1.2.0 is being used in the production areas at BESSY and SLAC with good results (stable). Release 1.3.0 is being tested.

4 Plans *What is on the wish list?*

The Gateway should work as tranparently as possible: A CA client should see no difference between a direct connection and a connection that is routed through a Gateway. While some differences are unavoidable, others should be worked upon.

There could be more documentation on how to use the Gateway.

Gateway configuration and additional debugging could be accessible through CA (restricted using Access Security).

There might be a chance to get “Virtual Gateways” running, i.e. multiple gateways running on different ports that may be configured independently, which share the PV data and attribute cache and its PV connections.

Any suggestions are very welcome. Volunteers for Beta-Testing, too.

CA Gateway Update

- Overview — *What is it? How does it work?*
- Applications — *What is it good for?*
- Status — *Latest improvements and bugfixes*
- Plans — *What is on the wish list?*

Overview

CA Uses One TCP Connection per Client/Server Combination

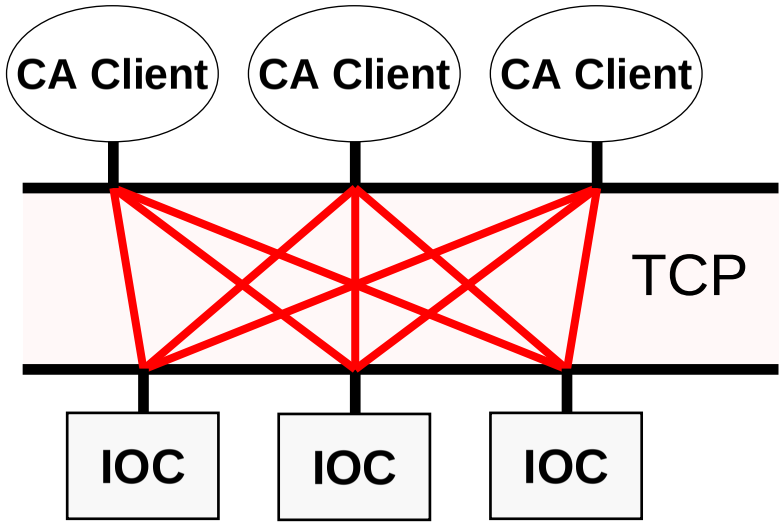

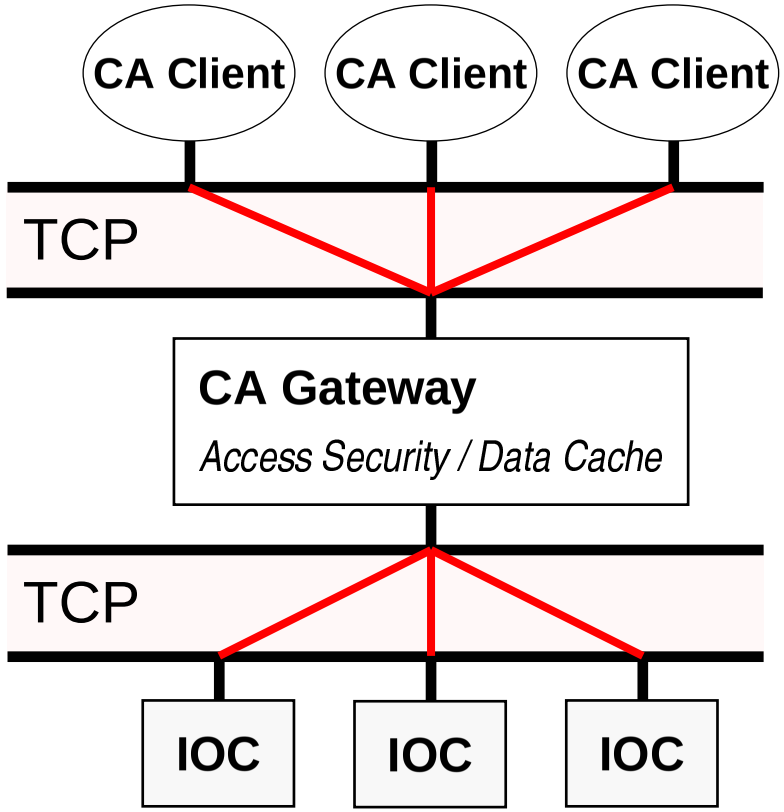

Overview

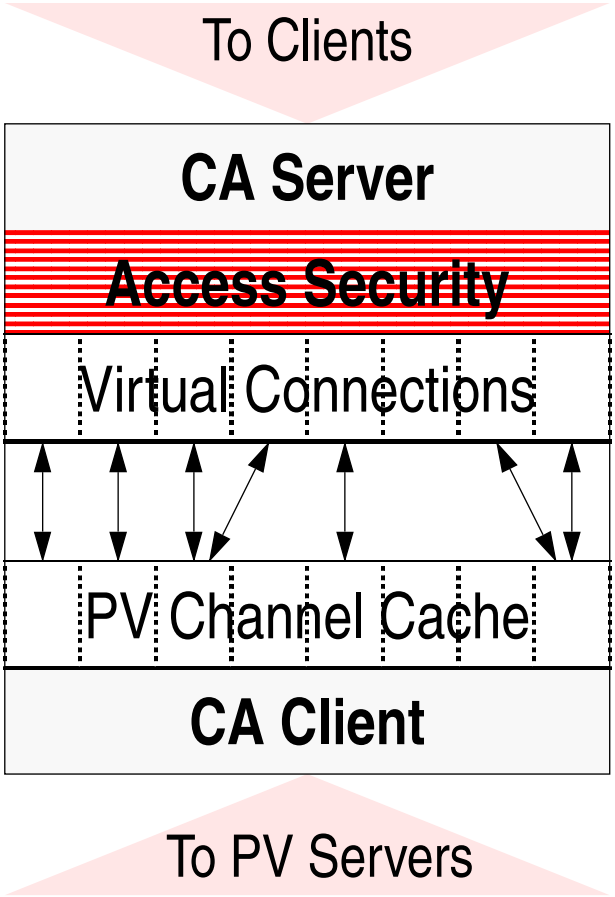

"Upper" Network

Portable CA Server

VC Names May Be "Real" PV Names or Aliases

Attributes and Values are Cached

"Lower" Network

- Reduces Number of TCP Connections
- Reads are Answered from Cache
- Events are Distributed to Clients
- PV Connections are Held Open

Applications

Controlled Access to Secure Nets

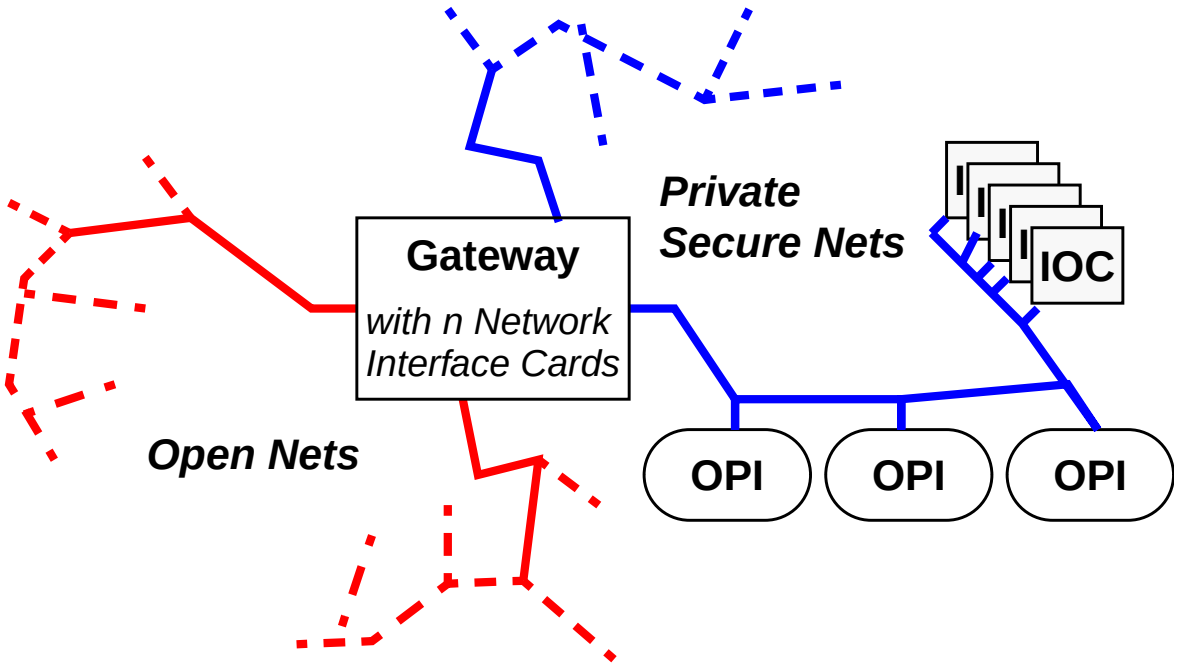

Channel Name Aliasing

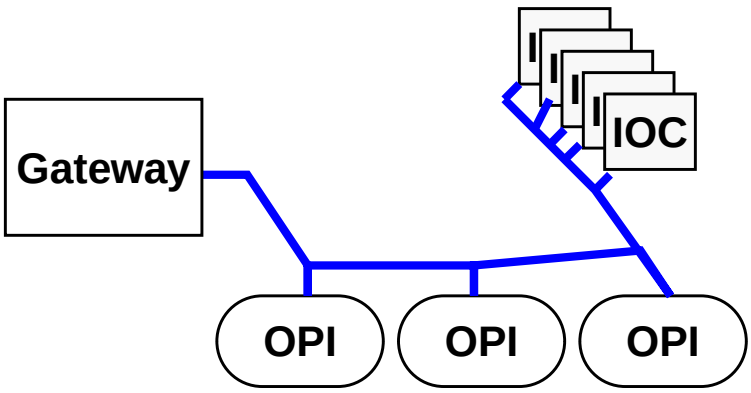

Applications

Panel Speed Up

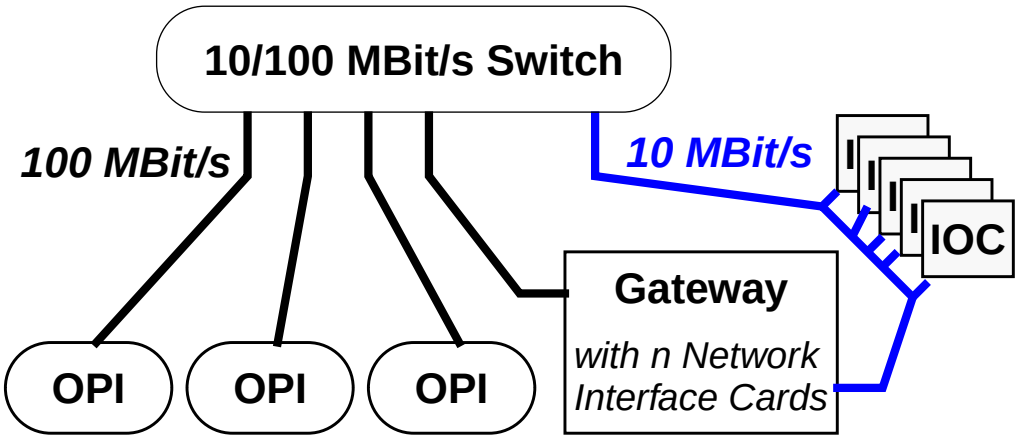

Local PV Repeater

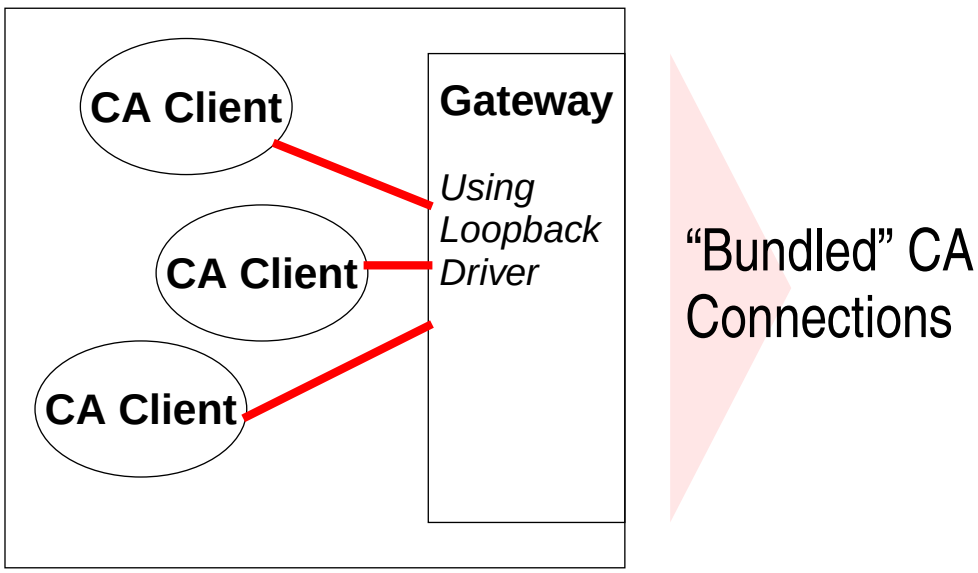

Status

- Alarm Handler Connections are Handled Properly
- Channel Deny may be Configured by Host Name
- Configuration Name Patterns use Regular Expressions (GNU Regex Library)
- IOC Reconnect Situations are Handled Correctly
- Still Needs a Patched Version of EPICS Base
- Used in Production System at APS and BESSY

Plans

- Increase Transparency:
  There Should be No Difference Between Accessing a PV Directly and via the Gateway
- More Documentation
- Configuration through Channel Access
- “Virtual Gateways”:
  Multiple Gateways Share One PV Data Cache
- Any Suggestions?