

EXPERIENCE WITH THE EPICS PV GATEWAY AT THE APS*

Kenneth Evans, Martin Smith, ANL, Argonne, IL 60439, U.S.A.

Abstract

The EPICS Process Variable Gateway [1] has become a stable, high-performance application that provides access to process variables while minimizing the impact on critical Input-Output Controllers (IOCs) and implementing additional access security. The additional access security typically prevents write access but is highly configurable. The Advanced Photon Source (APS) currently uses 40 Gateways running on 11 workstations to provide access to the machine subnet from the offices and for the individual experimental teams. These include reverse Gateways that allow administration of all 40 APS Gateways from a single MEDM [2] screen, even though the Gateways are running on separate networks. This administration includes starting, stopping, making and viewing reports, and viewing and editing access security files. There is one Gateway that provides process variable renaming. This paper provides an overview of the Gateways at the APS and describes the procedures that have been set up to use and administer them.

EPICS PV GATEWAY

The EPICS Process Variable Gateway is both a Channel Access server (like an IOC) and client (like MEDM). Clients such as MEDM connect to the server side, and the client side connects to IOCs and other servers. This allows many clients to access a process variable while making only one connection to the remote server, thus reducing the load on critical IOCs or other servers. In addition, it provides convenient access from one subnet to another and provides extensive additional access security. It also has the capability to provide aliases or renaming for process variables.

In addition to providing access to process variables in other servers, the Gateway publishes its own process variables, which allow it to be controlled and monitored using EPICS tools, such as MEDM and StripTool [3].

The Gateway is in use at many EPICS facilities and has become reasonably robust and powerful, handling thousands of process variables and high event rates, as described below. It runs on Solaris, Linux, other versions of UNIX, and Microsoft Windows and has been extensively used and tested on Solaris and Linux. The APS runs both Solaris and Linux Gateways.

APS GATEWAYS

Requirements

The APS has a single, machine subnet, where most of the IOCs and other servers are located. In addition, there are several subnets serving office areas used by APS

personnel, including scientists and engineers. There are experimental beamline groups on subnets largely administered by themselves that need access to machine process variables as well as to their own process variables.

The requirements are to provide read access to the machine process variables for the APS personnel and the experimental groups and in some cases write access for specified users as well.

There is currently a further requirement to provide an aliasing or renaming capability for a particular set of process variables, so they may be accessed by either of two sets of names. There is a requirement to provide access to video equipment, which is kept on a separate subnet than the machine subnet. Finally, there is a requirement to make the internal process variables of the Gateways themselves available so that they can be managed from a central location as described below.

Implementation

A diagram of the Gateways at the APS is shown in Fig. 1. The APS personnel are on Subnets 1 and 2 in the diagram and on other similar subnets. The video IOCs are on Subnet 2, which is also used by personnel. The experimental-group Gateways are set up in groups of four, running on one workstation along with a single reverse Gateway that has access to only the internal process variables of the other four plus its own. The grouping is similar to the way the experimental offices are grouped.

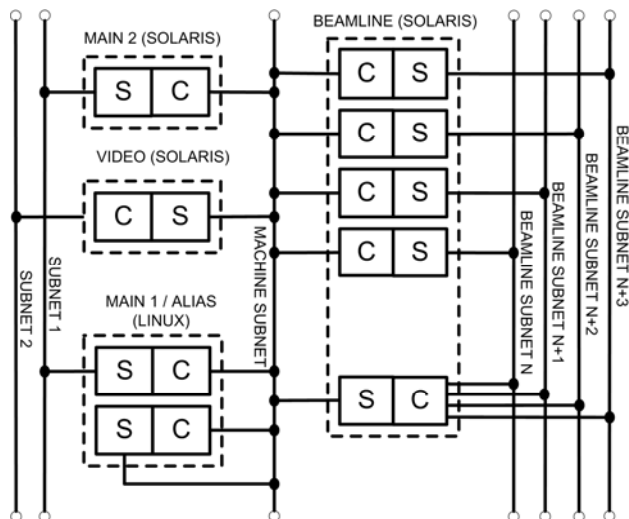


Figure 1: A diagram of the APS Gateways. The client side of a Gateway is denoted by C and the server side by S. There are eight groups of beamline Gateways, similar to the one shown, each serving four experimental groups and including a reverse Gateway to monitor the other four. The dotted lines represent a workstation.

*Work supported by the U. S. Department of Energy, Office of Basic Energy Sciences, under Contract No. W-31-109-ENG-38.

Only one of eight groups is shown in the diagram. Some of these groups have fewer than four Gateways, so there are a total of 28 Gateways plus eight reverse Gateways in the groups.

There are two main Gateways, which connect the machine subnet with office subnets like Subnet 1 and Subnet 2. The Main 1 Gateway has the most process variables and traffic of all of the Gateways and is running on Linux. All of the other Gateways are currently running on Solaris. The Main 2 Gateway is primarily used by a small group of floor coordinators and serves as a backup to the Main 1 Gateway.

The Video Gateway runs on its own workstation.

The Alias Gateway has both its client and its server sides on the machine subnet. The reason for this Gateway is to facilitate a large-scale renaming of process variables by doing it in stages. It only allows the specific process variables involved in this renaming. If a client accesses the name that is directly in the IOC, it gets it from the IOC. If it accesses the alias name, then it gets it from the Gateway. The Gateway merely uses its aliasing capability to provide access to the process variable in the IOC via a different name.

Security

The Gateway provides security and access control using regular expressions for process variable names and specification by user and host to either deny or allow access to these process-variable patterns. In order to facilitate the access specification, the APS has designated process variable names starting with upper-case letters to belong to the machine and those starting with lower-case letters or numbers for others.

MONITORING AND CONTROL

The APS Gateways are running on several different workstations on different subnets and do not have common file systems. Logging into each of them to start, stop, monitor, and make changes is inconvenient. It is possible, however, to monitor all of them from a single MEDM screen, shown in Fig. 2, and they are controlled by another screen, available through the "Gateway Controls" Related Display at the upper left in Fig. 2.

The reason the Gateways can be monitored and controlled by MEDM screens is that the Gateway has internal process variables that describe its state and others that can be used to make reports, reread the access

Sector:	Connected:	Active:	Inactive:	Dead:	Totals		Client Rates		Exist	Loop	CPU	Server Rates		Machine:
					VC:	PV:	Event:	Post:	Rate:	Rate:	Load:	Post:	Event:	
1	103	103	0	0	103	103	13.60	11.70	2.70	56.80	0.00	35.90	35.90	
2	168	168	0	0	168	168	54.10	53.90	2.40	67.40	0.00	456.70	456.70	gateway431
3	82	82	0	2	82	84	37.70	37.70	0.00	63.20	0.00	96.00	96.00	
4	128	128	0	0	128	128	28.60	24.60	0.00	58.20	0.01	402.09	402.09	
5	42	37	5	0	37	42	12.30	0.00	0.00	84.60	0.01	0.70	0.70	
6	33	33	0	0	33	33	2.00	2.00	0.00	54.70	0.00	10.70	10.70	gateway432
7	63	63	0	0	63	63	9.80	9.80	0.00	56.90	0.00	35.40	35.40	
8-BM	0	0	0	0	0	0	0.00	0.00	0.00	53.80	0.00	0.70	0.70	
9	79	79	0	0	79	79	23.00	23.00	61.40	62.00	0.01	29.70	29.70	
10	25	25	0	0	25	25	16.00	14.10	0.00	61.20	0.00	14.80	14.80	gateway433
11	231	223	8	0	223	231	78.70	3.70	0.00	59.80	0.01	4.40	4.40	
12	232	215	17	0	215	232	41.00	5.00	8.10	84.69	0.01	16.10	16.10	
13	33	33	0	0	33	33	3.90	3.90	0.80	55.30	0.00	29.30	29.30	
14	276	276	0	1	276	277	105.60	105.60	0.00	67.50	0.01	265.90	265.90	gateway434
15	58	58	0	0	58	58	12.20	12.20	0.00	56.20	0.00	22.90	22.90	
16	17	17	0	0	17	17	0.00	0.00	0.00	53.80	0.00	0.70	0.70	
17	10	7	3	0	7	10	3.90	2.00	0.00	62.70	0.00	2.70	2.70	
18	42	42	0	0	42	42	2.00	0.00	0.20	53.50	0.00	0.70	0.70	
19	29	29	0	0	29	29	4.00	4.00	25.00	54.80	0.00	10.40	10.40	gateway435
20	37	37	0	0	37	37	2.00	2.00	90.00	59.20	0.00	6.70	6.70	
21	0	0	0	0	0	0	0.00	0.00	0.00	54.10	0.00	0.70	0.70	
22	9	9	0	0	9	9	7.00	0.00	0.00	94.00	0.01	0.70	0.70	gateway436
23	28	28	0	0	28	28	9.30	9.30	0.00	59.10	0.00	28.60	28.60	
24	2	2	0	0	2	2	7.70	0.00	0.00	58.10	0.00	0.70	0.70	
31	41	41	0	0	41	41	12.50	3.90	0.00	62.20	0.00	5.60	5.60	
32	30	30	0	0	30	30	28.00	18.40	0.00	62.70	0.00	34.50	34.50	gateway438
33 & 34	200	40	160	0	40	200	12.80	8.60	27.60	97.69	0.03	86.70	86.70	
8-ID	29	29	0	0	29	29	3.10	3.10	0.00	55.70	0.00	20.70	20.70	gateway437
Main 1	16398	14033	2325	31	14033	16389	2259.48	2254.68	13.79	266.03	0.00	3091.63	3090.63	Ct.lapps41
Main 2														Rhea
r431	52	52	0	0	52	52	2.80	2.80	23.60	61.50	0.00	8.50	8.50	gateway431
r432	52	52	0	0	52	52	2.80	2.80	23.80	61.80	0.00	8.50	8.50	gateway432
r433	52	52	0	0	52	52	2.80	2.80	22.40	61.10	0.00	8.50	8.50	gateway433
r434	52	52	0	0	52	52	2.80	2.80	24.20	60.90	0.00	8.50	8.50	gateway434
r435	52	52	0	0	52	52	2.80	2.80	17.10	60.60	0.05	8.50	8.50	gateway435
r436	52	52	0	0	52	52	2.80	2.80	16.90	61.00	0.00	8.50	8.50	gateway436
r437	496	496	0	0	496	496	21.30	0.70	22.00	58.90	0.00	3.40	3.40	gateway437
r438	39	39	0	0	39	39	2.10	2.10	21.90	61.60	0.00	6.80	6.80	gateway438
rctlapps1	4	4	0	0	4	4	0.00	0.00	23.40	58.20	0.00	1.60	1.60	ctl.lapps1
alias	288	288	0	0	288	288	29.27	3.20	23.78	115.28	0.00	7.99	7.99	ctl.lapps41

Figure 2: MEDM screen that monitors and controls all the APS Gateways. The white boxes, indicating no connections, arise because the MEDM running this screen is attached to the Main 1 Gateway and cannot see the internal process variables of Main 2. This could be corrected, if desired, by another reverse Gateway.

security files, and stop it. Such internal statistics as the number of connected, active, inactive, and dead process variables; the number of virtual connections (clients such as MEDM on the server side); and the number of client process variables (from servers on the client side) are shown on the screen in Fig. 2, along with several event rates and the CPU usage.

The MEDM control screen has items to set the internal process variables that stop the Gateways, reread the access security, and make various reports. The Gateways are started by an MEDM Shell Command that runs a remote startup script using Secure Shell (SSH). This requires that the appropriate keys be previously placed on the remote workstations. Other MEDM Shell Commands are used to view the reports and view and edit access security files.

Reverse Gateways

The Gateway internal process variables are only published on the server side of the respective Gateways, and are not directly seen on the machine subnet. Moreover, the server sides are on several different subnets. For these process variables to be seen on a common MEDM screen, reverse Gateways are necessary. The reverse Gateways attach on their client side to the internal process variables of the regular Gateways. Since these reverse Gateways have their server side on the machine subnet, these process variables are then visible to all clients on the machine subnet. In order to prevent loops, the reverse Gateways are set to only attach to the internal process variables of themselves and the four regular Gateways, and the regular Gateways are set to deny the process variables from the other three regular Gateways.

The reverse Gateways see the same search request packets as the other servers on the machine subnet. One of the internal process variables they publish is their exist test rate, which is the same as the search request rate. Consequently, these Gateways are used to monitor the number of search requests happening on the machine

subnet. This information is archived and used in diagnosing network problems.

OPERATIONAL EXPERIENCE

The Main 1 Gateway, running on Linux, is the most used of the APS Gateways. It typically handles on the order of 10,000 – 20,000 process variables on its client side. Its server event rates are typically on the order of 2 - 10 kHz. It does this using a moderately small amount of CPU, on the order of 10% on a reasonably powerful, dual 2.8-GHz processor, Linux workstation. It has run as long as three months without having to be restarted. Typically it is stopped and restarted to upgrade to a new software version or because the workstation is rebooted for other reasons, not because of failures. It serves as a test bed for development and actually runs with debug, rather than optimized code. It has enough CPU headroom to do this.

CONCLUSIONS

The current version of the Gateway has been found to be quite stable and capable of handling a large number of process variables along with high update rates. It has run for long periods of time in a production environment without having to be restarted. It has become an important part of the APS Control System and is relied upon by many users. The entire system of APS Gateways can be monitored and controlled from a small number of MEDM screens. Reverse Gateways are used to make this possible.

REFERENCES

- [1] Gateway: The Process Variable Gateway,
<http://www.aps.anl.gov/epics/extensions/gateway/index.php>.
- [2] MEDM: Motif Editor and Display Manager,
<http://www.aps.anl.gov/epics/extensions/medm/index.php>.
- [3] StripTool,
<http://www.aps.anl.gov/epics/extensions/StripTool/index.php>.