

APS Controls Infrastructure Monitoring System Software Developers Guide for the EPICS Community

Deborah Quock

Revision Date: September 27, 2010

INTRODUCTION

An Infrastructure Monitoring System (IMS) has been developed at Argonne National Laboratory's Advanced Photon Source (APS) to provide immediate notification to the Controls Group on-call staff in the event of failure of critical controls hardware components or software processes. Approximately 3,000 operation-critical controls components and processes have been identified, many of which can be readily monitored through EPICS records. Nagios [1], an open-source host, service, and network monitoring program, has been interfaced to EPICS Channel Access monitoring tools to provide historical tracking of controls infrastructure events, email, and pager notifications to on-call staff, as well as Internet-accessible status displays. Two methods have been evaluated for checking EPICS alarm status data: a modified version of Mauro Giacchini's NAL [2] plug-in for Nagios, and a Perl script that uses a Channel Access Perl library. Personal Device Assistant-viewable Web pages are generated using data retrieved from the Nagios relational database (NDOUtils); along with data from the APS controls database IRMIS. Performance studies have validated the reliability of Nagios and the Perl Channel Access software.

A modern control system for a large accelerator consists of *thousands* of components, processes, and applications that must all be continuously functioning for optimum performance. A failure of one of these components may not be immediately noticeable by the machine operators and may continue for some time, possibly discovered several days later. The Infrastructure Monitoring System is intended to *exhaustively* monitor all parts of the control system and provide immediate notification to the on-call controls staff of an exception, in many times even before the machine operators notice the impact on machine performance. To meet these goals, the following project development steps were taken:

- Identify APS Controls IMS categories that will aid on-call staff in promptly deciphering cascading failures;
- Build a high data throughput, EPICS-based IMS software application; and
- Create PDA-friendly user displays of condensed IMS status reports.

APS CONTROLS IMS CATEGORIES

Nine major control system IMS categories were identified for creating a highly organized list of APS controls EPICS processes, controls hardware components, and other miscellaneous controls processes. These major IMS categories are:

- CCMS (Component Communication Monitoring System)
- Controls Servers
- Event Receivers
- IOCs
- Machine Status Links
- Nagios Software Architecture Components
- PV Gateway
- Timing
- VME/VXI Power Supplies

The major categories were then broken down into sub-categories to provide a three-level controls IMS grouping hierarchy with the lowest level being the actual EPICS process variable, server, or other type of computer process whose operating status is actually being checked and reported upon. For example, the major IMS category IOC has been divided into the following eight sub-categories:

- CALinks
- CPULoad
- LogServer
- Memory
- SaveRestore
- SoftHeartbeat
- TCPConnections
- TaskStatus

The subcategory CALinks contains approximately 800 EPICS process variable (PV) names where each PV name corresponds to a specific IOC's channel access connection status.

All IMS categories (with the exception of the IOC sub-category SaveRestore) have at this time been populated with either EPICS PVs, computer server names, or computer process names. Scripts written in PHP and Perl languages parse information out of the APS controls relational database IRMIS to generate lists of unique names for IMS.

APS CONTROLS IMS SOFTWARE ARCHITECTURE

Nagios software can monitor network services (e.g., SMTP, POP3, HTTP, PING) and host resources (e.g., processor load, disk usage). It also provides a simple plug-in design that allows users to develop their own service checks for unique situations. Starting with Mauro Giacchini's work at INFN on adapting EPICS channel access PV status updates to Nagios software as a plug-in, a full-scale, reliable controls IMS for APS was built with Nagios software as the foundation.

The Nagios definition of the states that a service or host can be in matches well with the EPICS PV alarm severity states shown in Table 1.

Table 1: EPICS-to-Nagios Translation

EPICS Alarm Severity	Nagios State
(none)	OK
Minor	Warning
Major	Critical
PV Readback Invalid	Critical
PV Name Not Found	Unknown

The Nagios feature to define network host hierarchy using host groups, hosts, and services was utilized when defining the APS Controls IMS categories in Nagios configuration files as shown here:

- Host Group
 - Host
 - Service
 - Service
 - ...
- IOC
 - CALinks
 - iocacis:ascaDisco
 - iocacis:dbcaDisco
 - ...

- CPUload
 - iocacis:load
 - iocbbpm1:load
 - ...

Figure 1 below shows APS Controls IMS categories as translated and displayed in the Nagios software “Hostgroup Summary.” Figures 2 and 3 provide an example of how a user can drill down to more detailed levels of information on each IMS Host Group category.




Figure 1: APS controls Nagios display.

Nagios - Windows Internet Explorer

http://orion/nagios/

Nagios

Nagios

General

- Home
- Documentation

Monitoring

- Tactical Overview
- Service Detail
- Host Detail
- Hostgroup Overview
- Hostgroup Summary
- Hostgroup Grid
- Servicegroup Overview
- Servicegroup Summary
- Servicegroup Grid
- Status Map
- 3-D Status Map
- Service Problems
- Unhandled
- Host Problems
- Unhandled
- Network Outages

Show Host:

- Comments
- Downtime
- Process Info
- Performance Info
- Scheduling Queue

Reporting

- Trends
- Availability
- Alert Histogram
- Alert History
- Alert Summary
- Notifications
- Event Log

Configuration

- View Config

Current Network Status

Last Updated: Fri Feb 13 16:06:45 CST 2009
 Updated every 90 seconds
 Nagios® 3.0.1 - www.nagios.org
 Logged in as *nagiosadmin*
 - Notifications are disabled

[View Status Overview For All Host Groups](#)
[View Service Status Detail For This Host Group](#)
[View Host Status Detail For This Host Group](#)
[View Status Summary For This Host Group](#)
[View Status Grid For This Host Group](#)

Up	Down	Unreachable	Pending
0	0	0	6
All Problems		All Types	
0		6	

OK	Warning	Unknown	Critical	Pending
71	0	0	0	0
All Problems		All Types		
0		71		

Service Overview For Host Group
'EVENT-RECEIVERS_HOSTGROUP'

[event-receivers_host_group \(EVENT-RECEIVERS_HOSTGROUP\)](#)

Host	Status	Services	Actions
EVENT-RECEIVERS_BOOSTER	PENDING	11 OK	
EVENT-RECEIVERS_Linac	PENDING	13 OK	
EVENT-RECEIVERS_PAR	PENDING	4 OK	
EVENT-RECEIVERS_RF	PENDING	17 OK	
EVENT-RECEIVERS_SR	PENDING	24 OK	
EVENT-RECEIVERS_Timing	PENDING	2 OK	

Figure 2: APS controls event-receivers Nagios display.




Figure 3: APS controls event-receivers timing Nagios display.

NAGIOS CONFIGURATION FILES AND APS CONTROLS IMS CUSTOM SCRIPTS

Nagios version 3.0.6 is currently installed and running on a Controls Group Linux server at APS. Download, installation, configuration and user manual information for Nagios is available at <http://www.nagios.org/>

The Nagios software configuration files comprise definitions of services and hosts that are to be monitored along with parameters for tuning the performance and logging of Nagios events. In-house written custom PHP and Perl scripts, and XML data files, have been created to generate the lists of EPICS PVs, Controls Group servers, and controls processes that are to be monitored by IMS. The APS Integrated Relational Model of Installed Systems (IRMIS) MySQL database is queried by the PHP and Perl scripts to generate lists of IMS EPICS PVs. Figure 4 describes the overall IMS process data flow.




Figure 4: IMS process data flow diagram.

The files for this application are located in an APS Controls Group CVS repository. The IMS CVS directory structure and file names are shown below.

```
cvsroot/infrastructmon/nagios/APS/  
    /CVS/  
    /etc/
```

```
cvsroot/infrastructmon/nagios/APS/Blowfish_PP.pm  
    /ccms_iocs.cfg  
    /ccms_pv_names.xml  
    /ccms_pv_search.php  
    /ccms_services.cfg  
    /CVS/  
    /eventreceivers_pv_names.xml  
    /generate_ims_controls_group_server_hosts@  
    /generate_ims_ioc_services@  
    /generate_ims_services@  
    /ims_controls_group_server_hosts.cfg  
    /ims_hosts.cfg  
    /ims_ioc_hosts.cfg  
    /ims_ioc_services.cfg  
    /ims_services.cfg  
    /ioc_pv_names.xml  
    /lib/  
    /logs/  
    /machinestatuslink_pv_names.xml  
    /nagioscamon@  
    /nagios_daily_reminder.pl  
    /nagios_db.properties  
    /pretend_nagios_pipe.txt  
    /pvgateway_pv_names.xml  
    /run_nagios_daily_reminder.sh  
    /test_pvgateway_with_nameserver.txt  
    /timing_pv_names.xml  
    /vmevxisupply_pv_names.xml
```

```
cvsroot/infrastructmon/nagios/etc/cgi.cfg  
    /CVS/  
    /mrtg/  
    /nagios.cfg  
    /ndo2db.cfg  
    /ndomod.cfg  
    /objects/  
    /resource.cfg
```

```
cvroot/infrastructmon/nagios/etc/objects/commands.cfg
    /contacts.cfg
    /CVS/
    /epicsExample.cfg
    /epics_hosts.cfg
    /ims_controls_group_server_hosts.cfg
    /ims_hostgroups.cfg
    /ims_hosts.cfg
    /ims_ioc_hosts.cfg
    /ims_ioc_services.cfg
    /ims_logserver_hosts.cfg
    /ims_logserver_services.cfg
    /ims_services.cfg
    /localhost.cfg
    /printer.cfg
    /switch.cfg
    /templates.cfg
    /timeperiods.cfg
    /windows.cfg
```

Many other Nagios files that come as part of the Nagios software installation package have not been included in the Controls Group CVS repository, but are currently installed and located under the directories:

```
usr/nagios/bin/
    /etc/
    /libexec/
    /sbin/
    /share/
    /var/
```

(Exception: `usr/nagios/etc/objects` directory is checked in/out from the Controls Group CVS repository.)

Almost all of the directories, subdirectories and files located in `usr/nagios` are owned by user “nagios” and group “nagios”. A few are currently shown as owned by user “root” and group “root” – these could probably be changed to user and group “nagios”.

NAGIOS INSTALLATION CONFIGURATION FILES

The main Nagios configuration file that comes as part of the Nagios software installation is called “nagios.cfg”. This file includes statements that define the directory location and names of files accessed by the Nagios process. This configuration file can be thought of as the main supervisory file that dictates where all other Nagios configuration files can be found. File `nagios.cfg` also includes directive type of statements on what operating parameters Nagios should run under. For example what level of error logging to perform, whether flapping detection should be initiated for all services and hosts upon startup, whether automated notifications such as email and pagers should be initiated for all services and hosts upon startup, etc. Caution should be exercised when modifying this configuration file as extensive testing has been undertaken at APS to optimize the performance of Nagios for monitoring several thousand services (here, EPICS PVs) and continuously writing this data to the Nagios MySQL database.

The Nagios users’ manual available at <http://www.nagios.org/> is the best reference source to learn about the numerous options for configuring Nagios. Several books have been published on Nagios that are also helpful. A recommended book is “Building a Monitoring Infrastructure with Nagios” by David Josephsen, ISBN 0132236931, Prentice Hall, copyright 2008.

COMPONENT COMMUNICATION MONITORING SYSTEM (CCMS) EPICS PVS

One of the goals in establishing a controls infrastructure monitoring system at APS was to include an I/O type of heartbeat EPICS PV for every hardware component that the Controls Group is responsible for and that communicates through an EPICS IOC. A PHP script was created to crawl the IRMIS database for components that have been pre-designated with the function “CCMS”. This script has file name `ccms_pv_search.php`. The IRMIS database schema for tables `component_type_function` and `function` is shown below in Figure 5. This script generated with a fair amount of accuracy EPICS PVs for approximately 25% of the known CCMS components. The remaining CCMS components were manually assigned a corresponding PV name to use in IMS by Gary Gunderson and were reviewed by the corresponding Controls Group staff control system cognizants. A total of 920 CCMS EPICS PVs are currently monitored in IMS and they are defined in file “`ccms_pv_names.xml`”.




Figure 5: IRMIS component type function database schema.

PERL SCRIPTS FOR GENERATING EPICS PVS FOR OTHER IMS CATEGORIES

There are three custom written Perl scripts used for generating APS XML data files (files contain lists of IMS PVs and are read by Perl script “`nagioscamon@`”) and APS specific Nagios configuration files (files read by the Nagios process when Nagios is started). These three Perl scripts are:

- `generate_ims_ioc_services@`
- `generate_ims_services@`
- `generate_ims_controls_group_server_hosts@`

As a general rule for the APS Controls IMS application, configuration files read by Nagios end with extension “.cfg” and files read by Perl script nagioscamon@ end with extension “.xml”. These generated configuration files are written in a format that is understood by the Nagios software: PVs are defined as being “services” and IMS categories are defined as being “hosts”. CCMS IRMIS component IDs are defined as being “service groups”. The top level of APS IMS categories are defined as “host groups”.

IOCs

The list of EPICS PVs for APS IOCs is generated by the Perl script “generate_ims_ioc_services@”. This script queries the IRMIS database for IOCs that have a status of ‘production’ or ‘ancillary’, and then generates the same group of PVs for all IOCs that is shown below:

- iocname:ascaDisco
- iocname:dbcaDisco
- iocname:seqcaDisco
- iocname:load
- iocname:memoryFree
- iocname:tcpConnections
- iocname:status

The exceptions to this general rule is that soft IOCs are assigned only one PV and that PV is iocname:heartbeat. Other exceptions to IOC IMS PVs are found in the Perl script “generate_ims_ioc_services@”. Each time this script is run, it generates the following files:

- ims_ioc_hosts.cfg
- ims_ioc_services.cfg
- ioc_pv_names.xml

Controls Group Servers

The file “generate_ims_controls_group_server_hosts@” is a Perl script that queries the table ‘server’ in the IRMIS database to retrieve a list of server names, and then generates the necessary Nagios configuration file “ims_controls_group_server_hosts.cfg”.

All Other IMS Categories

The file “generate_ims_services@” is a Perl script that reads all files with names ending in the string “_pv_names.xml” located in the same directory level and then generates configuration files “ims_services.cfg” and “ims_hosts.cfg”.

Caution: Remove or rename files ims_services.cfg and ims_hosts.cfg before running the Perl script generate_ims_services@. This is necessary because the script has been designed to repeatedly extract data from each XML file that it parses and then append this information to the host and services .cfg files.

The IMS XML files contain lists of EPICS PVs and at the time of this writing include:

- ccms_pv_names.xml
- eventreceivers_pv_names.xml
- ioc_pv_names.xml
- machinestatuslink_pv_names.xml
- pvgateway_pv_names.xml
- timing_pv_names.xml
- vmevxisupply_pv_names.xml

NAGIOS CHECK BY SSH COMMAND

The IMS subcategories IOC_LogServer and Nagios_MysqlServer monitor processes running on remote servers by using the Nagios check_by_ssh command.




Figure 6: IMS IOC_LogServer processes.

PERL SCRIPT FOR EPICS CHANNEL ACCESS MONITOR CALLBACKS

Currently, a mixture of approximately 3,000 EPICS PVs, computer servers, and computer processes are being monitored by Nagios software as Nagios “services.” All but 35 of these services are EPICS PVs. It is estimated that the number of EPICS PVs that will be included in the APS Controls IMS will ultimately grow to more than 5,000.

The EPICS plug-in for Nagios developed at INFN uses channel access “gets” and can be configured in an active polling mode within the Nagios software on an automated schedule. An example of such a Nagios configuration could be: once every three minutes, Nagios executes the EPICS plug-in for each EPICS PV that is listed as a “service” in Nagios. This type of active polling for such a large number of services in Nagios creates a burden on all IOCs visible to the plug-in, and Nagios software performance degradation is also noticeable. Consequently, an alternative to using the EPICS plug-in for checking on the status of an EPICS PV was developed to reduce both network traffic and processing of unnecessary information within Nagios.

Andrew Johnson of APS developed a Perl interface to the EPICS channel access library. The script “nagioscamon@” was written by Debby Quock specifically for the APS Controls IMS and it utilizes channel access monitor callbacks. This script on startup reads all files ending in text “_pv_names.xml” to extract the entire list of EPICS PV names, runs continuously, and *only* pushes “service status updates” to Nagios when a change in an EPICS PV alarm severity state is detected. Nagios has been configured at APS to accept *external passive checks* through a UNIX pipe for all of the Nagios services that are EPICS PVs. The server CPU usage of this Perl monitoring script for roughly 3,000 quiescent EPICS PVs is negligible (less than 1%). The UNIX pipe used by Nagios at APS is called “nagios.cmd” and is located in directory /usr/nagios/var/rw/

Caution: It is crucial that the directory “rw” is owned by user “nagios” in order for the Nagios Web page that is run on ctlappsirmis to have the authority to submit user requests initiated from the Nagios Web page to the Nagios process.

NAGIOS NDOUTILS DATABASE AND PDA-FRIENDLY WEB DISPLAYS

The Nagios software comes with many fine tools for such things as displaying alert history, producing availability reports, disabling and enabling configuration parameters. However, a phone call at two in the morning from an APS main control room operator to a soundly sleeping controls engineer can be startling. The engineer requires a simplified view of the controls problem at hand for prompt recovery action. This simplified view is created in PHP Web displays (Figures 6 through 8 below) that retrieve summary IMS status information from the Nagios NDOUtils MySQL relational database, and controls software and hardware “as-built” information from the APS controls relational database IRMIS. The Nagios MySQL database is installed on server ctrlrdbprod and the database name is “nagios”. The MySQL password for the nagios database can be obtained from the Controls Group leader. Configuration file ndo2db.cfg is installed on ctlappsirmis as part of the NDOUtils installation. More information about the Nagios NDOUtils relational database software can be found at <http://www.nagios.org/docs/>

Three levels of IMS PDA-friendly Internet displays are provided through the secure APS Web site <https://www.aps.anl.gov/controls/ims.php> and allow a user to drill down to more levels of detail as needed. The top level display in this directory is file “ims.php”. These files are currently *not* under CVS control.




Figure 7: IMS PDA-friendly top level overall status display.




Figure 8: IMS PDA-friendly second level, critical and warning services display.




Figure 9: IMS PDA-friendly third level, service details display.

NAGIOS DAILY REMINDER TO ON-CALL BLACKBERRY

The Perl script “nagios_daily_reminder.pl” was created to be executed once a day by a Cron job for the purpose of querying the Nagios database for the current overall status of IMS and then sending this information to the Controls Group on-call staff in an email. (File “run_nagios_daily_reminder.sh” is currently run as a scheduled Cron job on server ctappsirmis.) This email also serves the purpose of reminding the Controls Group on-call staff to attend the APS main control room daily 4 p.m. meeting. An example of the IMS daily reminder email message format is shown below. The number of critical services calculated by the Perl script and displayed in the email message does *not* include services that have a user entered Nagios comment associated with them.

```
Controls IMS Status Report at: Mon Feb 16 15:33:55 2009
```

```
IMS Service Status Totals
OK:      2945
WARNING:  1
UNKNOWN:  0
ACTIVE CRITICAL:  0
```

STARTING AND STOPPING PROCESSES NAGIOS, NDO2DB AND NAGIOSCAMON

Required EPICS environment variables for running Nagios on server ctappsirmis are:

```
EPICS_CA_AUTO_ADDR_LIST=YES
EPICS_CA_ADDR_LIST=###.##.###.## (PV gateway address)
EPICS_HOST_ARCH=linux-x86 (server ctappsirmis operating system)
```

There is a specific order in which the Nagios database utility ndo2db, Nagios process, and Perl script nagioscamon@ must be started and stopped.

The database utility ndo2db must be running before the Nagios process is started. To start the ndo2db process, use the line command:

```
bash$ /etc/init.d/ndo2db start
```

(will be prompted for Nagios password)

Next, start the Nagios process by typing:

```
bash$ sudo /sbin/service nagios start
```

(will be prompted for Nagios password)

The last step is to run the Perl script nagioscamon@ on server ctappsirmis. This is simply done by CD'ing to the directory /infrastructmon/nagios/APS and typing line command:

```
bash$ nagioscamon@
```

(will display repeated lines of output to the monitor terminal window as it runs continuously)

The correct order in which to stop above processes is as follows:

First: type CNTL-C to stop execution of nagioscamon@

Second: type line command

```
bash$ sudo /sbin/service nagios stop
```

(to cleanly kill the nagios process)

Third: type line command

```
bash$ /etc/init.d/ndo2db stop
```

(to cleanly kill the ndo2db process)

ADDING A COMMENT TO A NAGIOS SERVICE

A comment can be added to a Nagios service and saved in a persistent state. The example below shows a CCMS PV that is in a critical state (highlighted with red background). To navigate to this Nagios service (a.k.a. EPICS PV), simply click on the red colored link and a new Nagios display (Figure 10) will appear that has more detail on this particular service. By selecting the underlined service name (Web link) in Figure 10, another display will open (see Figure 11) that has service state information and a link near the bottom of the display for adding or deleting comments for this specific service.

Note that the IMS PDA-friendly PHP source code and the IMS daily reminder Perl script have been written to take into account which IMS services have user defined comments associated with them. The Nagios NDOUtils MySQL database stores information on Nagios service comments, and this is in turn is queried by the IMS PDA-friendly PHP and IMS daily reminder Perl scripts for refining the information presented to Controls Group on-call staff.

Caution: The Nagios main configuration file nagios.cfg must have parameter setting “ retain_state_information=1 ” so that upon startup, Nagios will reload all saved service and host state information before starting to monitor. This includes knowledge of user defined comments for Nagios services.




Figure 10: Nagios overall status display with one critical service.

Nagios - Windows Internet Explorer

http://orion/nagios/

Nagios

Nagios

General

- Home
- Documentation

Monitoring

- Tactical Overview
- Service Detail
- Host Detail
- Hostgroup Overview
- Hostgroup Summary
- Hostgroup Grid
- Servicegroup Overview
- Servicegroup Summary
- Servicegroup Grid
- Status Map
- 3-D Status Map
- Service Problems
 - Unhandled
 - Host Problems
 - Unhandled
 - Network Outages
- Show Host:
- Comments
- Downtime
- Process Info
- Performance Info
- Scheduling Queue

Reporting

- Trends
- Availability
- Alert Histogram
- Alert History
- Alert Summary
- Notifications
- Event Log

Configuration

- View Config

Current Network Status
 Last Updated: Fri Feb 13 16:08:07 CST 2009
 Updated every 90 seconds
 Nagios® 3.0.1 - www.nagios.org
 Logged in as nagiosadmin
 - Notifications are disabled

[View Service Status Detail For All Host Groups](#)
[View Host Status Detail For This Host Group](#)
[View Status Overview For This Host Group](#)
[View Status Summary For This Host Group](#)
[View Status Grid For This Host Group](#)

Host Status Totals

Up	Down	Unreachable	Pending
0	0	0	104

[All Problems](#) [All Types](#)

0	104
---	-----

Service Status Totals

Ok	Warning	Unknown	Critical	Pending
918	0	0	1	0

[All Problems](#) [All Types](#)

1	920
---	-----

Display Filters:
 Host Status Types: All
 Host Properties: Any
 Service Status Types: Critical
 Service Properties: Any

Service Status Details For Host Group 'CCMS_HOSTGROUP'

Host	Service	Status	Last Check	Duration	Attempt	Status Information
CCMS_locs19vp	S20.A.Q2.MultAI	CRITICAL	02-13-2009 16:07:43	4d 0h 30m 54s	3/3	INVALID: S20:A:Q2.MultAI

1 Matching Service Entries Displayed

Figure 11: Nagios service display with comment shown as cloud icon.




Figure 12: Nagios display of service state information.

APS NAGIOS AUTHENTICATION AND PASSWORDS

The installation of Nagios software includes a password file called “htpasswd.users”. An explanation of how to create and modify this file is available in the Nagios installation guide. The main Nagios configuration file, /usr/nagios/etc/nagios.cfg, includes authorization parameters that need to set for defining authenticated users and the Nagios features they are allowed access to.

REFERENCES

- [1] E. Galstad, Nagios Enterprises, LLC, <http://www.nagios.org>.
- [2] M. Giacchini, Laboratori Nazionali di Legnaro, Istituto Nazionale di Fisica Nucleare, <http://www.lnl.infn.it/~epics/NAL.html>.