



View for ClusterStor™ Installation and Configuration Guide

(1.0.1)

S-3025

Contents

1 About View for ClusterStor™ Installation and Configuration Guide	3
2 Install View for ClusterStor™ Software.....	7
3 ClusterStor System Configuration for Metrics.....	9
4 ClusterStor System Configuration for Log Forwarding.....	11
5 Configuration for Lustre Job Statistics.....	13
5.1 Enable Lustre Job Statistics for CLE 5.2.....	13
5.2 Enable Lustre Job Statistics for CLE 6.0.....	14
5.3 Disable Lustre Jobstats for CLE 5.2.....	15
5.4 Disable Lustre Jobstats for CLE 6.0.....	16
5.5 Enable Jobstats on the ClusterStor System.....	16
5.6 Disable Jobstats on the ClusterStor System.....	17
6 ClusterStor System Configuration for SNMP.....	18
7 SMW Installation and Upgrade for Job Events.....	19
7.1 Configuration for CLE 5.2.....	19
7.2 Configuration for CLE 6.0.....	20
8 Add Initial ClusterStor to View for ClusterStor™	22
9 Configure Notification Email.....	23
10 Change Default Password.....	24
11 Update View for ClusterStor™ Software.....	25
12 Update Retention Policies for InfluxDB.....	26

1 About View for ClusterStor™ Installation and Configuration Guide

Scope and Audience

View for ClusterStor™ is a monitoring and metrics software package created by Cray which collects and persists performance and job metrics specific to the Cray ClusterStor storage system. View for ClusterStor™ collects Lustre® performance, jobs metrics, and system events specific to the storage system. Additionally, system logs, system metrics, and system events from each ClusterStor storage system can be configured to be monitored and will collect and persist ibstats metrics from the InfiniBand fabric if connected to the ClusterStor high speed InfiniBand network. View for ClusterStor™ can be integrated with the Cray System Management Workstation (SMW), and will collect job information such as start/stop, job id, ap id, user id, and duration for jobs launched on attached Cray computers. Administrators can view this information as it occurs or look at information collected at different points in the past through data dashboards and workflows. The *View for ClusterStor™ Installation and Configuration Guide (S-3025)* covers all of the necessary procedures for preparing an on-site server for initial View for ClusterStor™ use. This guide assumes the reader is familiar with the Cray ClusterStor storage system.

Release Information

The *View for ClusterStor™ Installation and Configuration Guide* supports View for ClusterStor™ version 1.0.1.

Product Requirements

Supported Versions

The following versions are supported by View for ClusterStor™:

- CentOS 7.2
- Software Docker CE 17.06.2 or greater
- Docker Compose 1.14.0 or greater

The ClusterStor Storage System must be running the following software releases or greater:

- ClusterStor 2.0 SU26
- ClusterStor 3.0 SU10

All Lustre Clients must be running the following release level and patch level or greater:

- CLE 5.2UP04 patch number PS281
- CLE 6.0UP02 patch number PS44
- CLE 6.0UP03 patch number PS15

The following ports are used by View for ClusterStor™:

- 80

- 441
- 9092
- 2128
- 8514

Hardware

- Standalone server with:
 - 128 GBytes of Memory or more
 - 8 CPU cores or more
 - SSD storage of 500GB or more

Time Synchronization

Time and time zone must be synchronized across any and all servers, especially:

- System Management Workstation (SMW)
- Storage server
- View for ClusterStor server
- Cray compute system

Hostname

The hostname of the View for ClusterStor™ server must match the DNS entry, if DNS is being used.

Record of Revision

Revision	Date	Content Information
<i>View for ClusterStor™ Installation and Configuration Guide (S-3025) 1.0.1</i>	05/16/2018	Release 1.0.1
<i>View for ClusterStor™ Installation and Configuration Guide (S-3025) Rev A</i>	03/20/2018	Revision A
<i>View for ClusterStor™ Installation and Configuration Guide (S-3025)</i>	03/14/2018	GA release

Typographic Conventions

<code>Monospace</code>	Indicates program code, reserved words, library functions, command-line prompts, screen output, file/path names, key strokes (e.g., <code>Enter</code> and <code>Alt-Ctrl-F</code>), and other software constructs.
Monospaced Bold	Indicates commands that must be entered on a command line or in response to an interactive prompt.
<i>Oblique or Italics</i>	Indicates user-supplied values in commands or syntax definitions.
Proportional Bold	Indicates a graphical user interface window or element.

\ (backslash) At the end of a command line, indicates the Linux® shell line continuation character (lines joined by a backslash are parsed as a single line). Do not type anything after the backslash or the continuation feature will not work correctly.

smaller font size Some screenshot and code examples require more characters than are able to fit on a line of a PDF file, resulting in the code wrapping to a new line. To prevent wrapping, some examples are displayed with a smaller font to preserve the file formatting.

Command Prompt Conventions

Host name and account in command prompts The host name in a command prompt indicates where the command must be run. The account that must run the command is also indicated in the prompt.

- The `root` or super-user account always has the `#` character at the end of the prompt.
- Any non-`root` account is indicated with `account@hostname>`. A user account that is neither `root` nor `crayadm` is referred to as `user`.

<code>smw#</code>	Run the command on the SMW as <code>root</code> .
<code>sdb#</code>	Run the command on the SDB node as <code>root</code> .
<code>boot#</code>	Run the command on the boot node as <code>root</code> .
<code>login#</code>	Run the command on any login node as <code>root</code> .
<code>hostname#</code>	Run the command on the View for ClusterStor™ system as <code>root</code> .
<code>user@hostname></code>	Run the command on the specified system as any non- <code>root</code> user.

Directory path in command prompt Example prompts do not include the directory path, because long paths can reduce the clarity of examples. Most of the time, the command can be executed from any directory. When it matters which directory the command is invoked within, the `cd` command is used to change into the directory, and the directory is referenced with a period (`.`) to indicate the current directory.

For example, here are actual prompts as they appear on the system:

```
smw:~ # cd /etc
smw:/etc# cd /var/tmp
smw:/var/tmp# ls ./file
smw:/var/tmp# su - crayadm
crayadm@smw:~> cd /usr/bin
crayadm@smw:/usr/bin> ./command
```

And here are the same prompts as they appear in this publication:

```
smw# cd /etc
smw# cd /var/tmp
```

```
smw# ls ./file  
smw# su - crayadm  
crayadm@smw> cd /usr/bin  
crayadm@smw> ./command
```

2 Install View for ClusterStor™ Software

Prerequisites

All site specific configuration values have been collected and the directory on the SSD for storing all SMA data has been created.

CentOS has been installed and the IB interface has been configured and is up and running.

Firewalld is not currently supported by View for ClusterStor™ and must be disabled before starting Docker. IPTables can be used to manage the firewall directly.

Docker 17.06.2-ce or greater has been installed from <https://docs.docker.com/install/linux/docker-ce/centos/>.

Docker Compose 1.14.0 or greater has been installed from <https://docs.docker.com/compose/install/>.

View for ClusterStor™ software package has been downloaded from CrayPort.

Procedure

1. Login to the View server as `root`.
2. Copy the View for ClusterStor™ software archive to `/root`.

The archive name will be a variation of `sma-1.0.0-<date-time stamp>.tgz`, e.g. `sma-1.0.0-201803121200.tgz`.

3. Unpack the archive using the tar command as shown below.

```
hostname# tar xvf sma-1.0.0-<date-time stamp>.tgz
```

4. Move to newly created directory, `/root/sma-install`.

```
hostname# cd /root/sma-install
```

5. Run the installation script.

If a default value is defined when prompted for site specific information, press **Enter** to use the default value.

Installation should take no more than forty minutes.

```
hostname# ./setup.sh
```

- a. Fill in the data directory path when prompted.

This is the directory on the SSD that will store all SMA data and configuration files.

```
Enter path for data directory (default= /var/sma/data):/site/data/directory
```

- b. Enter the FQDN hostname of the View for ClusterStor™ system.

```
Enter FQDN hostname (default= <detected hostname>): hostname.customer.site.com
```

- c. Enter the site email relay host.

```
Enter site email relay host: mail-relay.customer.site.com
```

- d. Change any previous answers if corrections need to be made.

Answer no if all config settings are correct.

```
Do you need to change any of the previous answers? (possible choices=[yes  
no]): no
```

6. Verify that the SMA service is running.

```
hostname# systemctl status sma
```

View for ClusterStor™ has now been installed.

3 ClusterStor System Configuration for Metrics

About this task

This step enables the View for ClusterStor™ system to distribute metrics. This SDK provides the ability to stream metric data off of a ClusterStor system and onto the SMA server. The following steps can be done on a “live” system, i.e. with Lustre targets mounted on the Lustre servers.

Procedure

1. Log in to the target ClusterStor management node (n000) as admin.
2. Change to the root user.

```
admin@MGMT0> sudo su -
password for admin: password
Last login: Thu Sep 29 12:47:14 CDT 2016 from 172.16.2.3 on ssh
```

3. Create a new user with read-only access.

There will be a prompt to create and confirm a password for the new user.

```
MGMT0# cscli admins add --username=smauser --role=readonly --disable-ssh --
enable-web
```

```
Enter the password : password
Confirm the password : password
```

```
MGMT0# cscli admins list
```

```
-----
Username  Role      Uid    SSH Enabled  Web Enabled  Policy
-----
smauser   readonly  1016   False        True         default
-----
```

4. Enable the REST API.

```
MGMT0# cscli service_console configure rest_api enable
```

5. Add the guest user as an REST API authorized user.

```
MGMT0# cscli service_console configure rest_api user_add --username smauser
User 'smauser' has been added to REST API authorized users list
```

6. Confirm the REST API.

```
MGMT0# cscli service_console configure rest_api show
REST API access: enabled
```

REST API authorized users:

```
smauser
```

7. Check if the system identifier is set. In this example, `snx11103n000` is used, along with its corresponding serial number.

```
MGMT0# cscli service_console configure system show
System settings:
  System serial number: CSSX0G4DE5
  System identifier name: [not-set]
```

- a. Assign the system identifier if it is not set.

```
MGMT0# cscli service_console \
configure system identifier -n snx11103n000
```

- b. Confirm the system identifier setting.

```
MGMT0# cscli service_console configure system show
System settings:
  System serial number: CSSX0G4DE5
  System identifier name: snx11103n000
```

The ClusterStor system has been configured for Metrics.



CAUTION: After the installation of a System Update (SU) on the ClusterStor system, the configuration will be cleared and will need to be performed again.

4 ClusterStor System Configuration for Log Forwarding

About this task

Designate the destination to which the ClusterStor system will send logs as well as which logs to send. For the sake of this example, SMA1234 is the given destination, and the View for ClusterStor™'s IP address is 172.30.76.13.

Procedure

1. Log in to both the active and standby target ClusterStor management nodes (n000 and n001) as admin.
2. Change to the root user.

```
admin@MGMT0> sudo su -
password for admin: password
Last login: Thu Sep 29 12:47:14 CDT 2016 from 172.16.2.3 on ssh
```

3. Save a copy of the `syslog-ng_receiver.erb` config file on both nodes.

```
MGMT0# cp /etc/puppet/modules/syslog_ng/templates/syslog-ng_receiver.erb \
/etc/puppet/modules/syslog_ng/templates/syslog-ng_receiver.save
```

4. Add lines to the `/etc/puppet/modules/syslog_ng/templates/syslog-ng_receiver.erb` file on both nodes.

```
MGMT0# vi /etc/puppet/modules/syslog_ng/templates/syslog-ng_receiver.erb
```

- a. Add before existing destination lines:

```
destination SMA1234 { udp("172.30.76.13" port(8514) time-zone("+00:00")); };
```

- b. Add before existing log lines:

```
log { source(s_sys); source(s_udp); destination(SMA1234); };
```

5. Run puppet to update the official files and integrate changes into the ClusterStor cluster on the active management node. The changes will appear in `/etc/syslog-ng/syslog-ng.conf`.

```
MGMT0# puppet agent -tv
```

6. Verify that the log forwarding changes have been successfully propagated on the active management node.

```
MGMT0# grep SMA1234 /etc/syslog-ng/syslog-ng.conf
```

The passive node will not update the `syslog-ng.conf` file until it is made the active node.

The logs have now been configured for forwarding on the ClusterStor system.

5 Configuration for Lustre Job Statistics

Prerequisites

The ClusterStor system has been configured for Metrics.

About this task

The default Lustre jobstats code on the client extracts the unique JobID from an environment variable within the user process, and sends this JobID to the server with the I/O operation. This environment variable lookup on the client causes Lustre I/O performance degradations when jobstats are enabled. A Cray Lustre client patch is required to workaround this performance issue and is available for CLE 5.2 versions UP04 and above, as well as CLE 6.0 versions UP02-UP04. With this patch, enabling Lustre jobstats will be done with a WLM prologue script when the job is launched.

The following procedure assumes that a prolog script and epilog script do not currently exist. Since ALPS only supports one prolog/epilog script, additional scripts, such as RUR scripts, need to be combined with the Lustre jobstats prolog/epilog scripts. For example, run the RUR prolog and epilog scripts inside the Lustre jobstats scripts.

Users running CLE 5.2 versions UP04 and above are required to perform the procedure in [Enable Lustre Job Statistics for CLE 5.2](#) on page 13 before proceeding. Users running CLE 6.0 versions UP02-UP04 are required to perform the procedure in [Enable Lustre Job Statistics for CLE 6.0](#) on page 14 before proceeding. Users running CLE 6.0 UP05 and above can proceed directly to [Enable Jobstats on the ClusterStor System](#) on page 16.

5.1 Enable Lustre Job Statistics for CLE 5.2

Procedure

1. Edit the `alps.conf` file on the shared-root file system on the boot node to set the prologue/epilogue paths.

```
boot# xtopview -m "add prolog/epilog scripts to alps.conf for Lustre nodelocal
jobstats"
boot# vi /etc/opt/cray/alps/alps.conf

prologPath      /ufs/alps_shared/lustre_jobstats_prolog.sh
epilogPath      /ufs/alps_shared/lustre_jobstats_epilog.sh
```

2. Create and install the `lustre_jobstats_prolog.sh` script on one of the login nodes to set "nodelocal" mode before the application starts.

```
login# vi /ufs/alps_shared/lustre_jobstats_prolog.sh

#!/bin/bash
export CRAY_ROOTFS=INITRAMFS
/opt/cray/nodehealth/default/bin/pcmd -f $ALPS_PREP_NIDFILE "/sbin/lctl
set_param jobid_var=nodelocal"
/opt/cray/nodehealth/default/bin/pcmd -f $ALPS_PREP_NIDFILE "/sbin/lctl
set_param jobid_name=$ALPS_PREP_APID"
```

3. Create and install the `lustre_jobstats_epilog.sh` script on one of the login nodes to disable Lustre jobstats when the application exits.

```
login# vi /ufs/alps_shared/lustre_jobstats_epilog.sh

#!/bin/bash
export CRAY_ROOTFS=INITRAMFS
/opt/cray/nodehealth/default/bin/pcmd -f $ALPS_PREP_NIDFILE "/sbin/lctl
set_param jobid_var=disable"
```

4. Set the correct permissions on the prolog and epilog scripts on one of the login nodes.

```
login# chmod 700 /ufs/alps_shared/lustre_jobstats_prolog.sh \
/ufs/alps_shared/lustre_jobstats_epilog.sh
```

5. Send a SIGHUP signal to `apsys` to reread the updated `apsys` configuration file.

This must be performed on each login node.

```
login# pkill -HUP apsys
```

5.2 Enable Lustre Job Statistics for CLE 6.0

Procedure

1. Invoke the configurator in interactive mode on the SMW for the `cray_alps` service to set the prologue/epilogue paths.

Replace `p0.staging` with the site-specific config set name.

```
smw# cfgset update -m interactive -s cray_alps p0.staging

cray_alps.settings.apsys.data.prologPath: /home/crayadm/bin/
lustre_jobstats_prolog.sh
cray_alps.settings.apsys.data.epilogPath: /home/crayadm/bin/
lustre_jobstats_epilog.sh
```

2. Run the configurator to confirm the settings.

```
smw# cfgset search -s cray_alps --level basic p0.staging

# 4 matches for '.' from cray_alps_config.yaml
#-----
cray_alps.settings.common.data.xthostname: crayxc
cray_alps.settings.common.data.alps_node_groups: [ ] # (empty)
```

```
cray_alps.settings.apsys.data.prologPath: /home/crayadm/bin/
lustre_jobstats_prolog.sh
cray_alps.settings.apsys.data.epilogPath: /home/crayadm/bin/
lustre_jobstats_epilog.sh
```

- Update the alps configuration set on the boot, sdb, and all login nodes.

```
boot# /etc/init.d/cray-ansible start
sdb# /etc/init.d/cray-ansible start
login# /etc/init.d/cray-ansible start
```

- Create and install the `lustre_jobstats_prolog.sh` script on one of the login nodes to set "nodelocal" mode before the application starts.

```
login# vi /home/crayadm/bin/lustre_jobstats_prolog.sh

#!/bin/bash
export CRAY_ROOTFS=INITRAMFS
/opt/cray/nodehealth/default/bin/pcmd -f $ALPS_PREP_NIDFILE "/sbin/lctl
set_param jobid_var=nodelocal"
/opt/cray/nodehealth/default/bin/pcmd -f $ALPS_PREP_NIDFILE "/sbin/lctl
set_param jobid_name=$ALPS_PREP_APID"
```

- Create and install the `lustre_jobstats_epilog.sh` script on one of the login nodes to disable Lustre jobstats when the application exits.

```
login# vi /home/crayadm/bin/lustre_jobstats_epilog.sh

#!/bin/bash
export CRAY_ROOTFS=INITRAMFS
/opt/cray/nodehealth/default/bin/pcmd -f $ALPS_PREP_NIDFILE "/sbin/lctl
set_param jobid_var=disable"
```

- Set correct permissions on the prolog and epilog scripts on one of the login nodes.

```
login# chmod 700 /home/crayadm/bin/lustre_jobstats_prolog.sh \
/home/crayadm/bin/lustre_jobstats_epilog.sh
```

- Restart the apsys daemon to reread the updated apsys configuration file.

This must be performed on all login nodes.

```
login# systemctl restart apsys
```

5.3 Disable Lustre Jobstats for CLE 5.2

About this task

In order to stop collecting jobstats, services must be disabled.

Procedure

1. Remove the prologue/epilogue paths from the ALPS configuration or disable the `lctl set_param` commands in the `prolog/epilog` scripts on the login node.
2. Send a `SIGHUP` signal to `apsys` to reread the update `apsys` configuration file.

This must be performed on all of the login nodes.

```
login# pkill -HUP apsys
```

5.4 Disable Lustre Jobstats for CLE 6.0

About this task

In order to stop collecting jobstats, services must be disabled.

Procedure

1. Remove the prologue/epilogue paths from the ALPS configuration or disable the `lctl set_param` commands in the `prolog/epilog` scripts on the login node.
2. Restart the `apsys` daemon to reread the updated `apsys` configuration file.

This must be performed on all login nodes.

```
login# systemctl restart apsys
```

5.5 Enable Jobstats on the ClusterStor System

Prerequisites

Either [Enable Lustre Job Statistics for CLE 5.2](#) on page 13 or [Enable Lustre Job Statistics for CLE 6.0](#) on page 14 has been completed.

Procedure

1. Enable Lustre job statistics collection from the ClusterStor management node.

```
MGMT0$ cscli lustre jobstats collection --enable
lustre: Enabling Lustre Job Statistics for snx
lustre: Updating puppet configuration. This can take a while...
lustre: Successfully enabled Lustre Job Statistics for snx.

MGMT0$ cscli lustre jobstats modify --frequency 30 --scheduler ALPS_APP_ID
lustre: Configuring Lustre Job Statistics for snx
```



```
lustre: Updating configuration. This can take a while...
lustre: Successfully configured Lustre Job Statistics for snx
```

2. Check the Lustre job statistics configuration to ensure it is enabled.

```
MGMT0$ cscli lustre jobstats list
-----
FSName      Collection  Frequency(in sec)  Scheduler
-----
snx Enabled      30                  ALPS_APP_ID
-----
```

5.6 Disable Jobstats on the ClusterStor System

Prerequisites

The procedure in [Disable Lustre Jobstats for CLE 5.2](#) on page 15 or [Disable Lustre Jobstats for CLE 6.0](#) on page 16 has been completed.

Procedure

1. Disable Lustre job statistics collection from the ClusterStor management node.

```
MGMT0$ cscli lustre jobstats collection --disable
```

2. Check the Lustre jobstats configuration to confirm that collection has been disabled.

```
MGMT0$ cscli lustre jobstats list
-----
FSName      Collection  Frequency(in sec)  Scheduler
-----
snx Disabled      30                  ALPS_APP_ID
-----
```

6 ClusterStor System Configuration for SNMP

About this task

Enabling Simple Network Management Protocol (SNMP) on ClusterStor permits View for ClusterStor™ to obtain information about service alerts and determine what alarm state a system may be in at any given time. The View for ClusterStor™ interface will display these service alerts, and the SNMPwalk interface will create and send an email with more detailed information.

Procedure

1. Log in to the target ClusterStor management node (n000) as admin.
2. Enable SNMP.

```
admin@MGMT0> sudo cscli service_console configure snmp enable
Attempting to enable SNMP, please wait
.....
SNMP has successfully been enabled.
```

SNMP is now enabled.

7 SMW Installation and Upgrade for Job Events

About this task

The jobevent daemon captures the WLM job and application start/stop events on the SMW and forwards them to the View server. Following are the required steps to install or upgrade the daemon. Configuration steps are dependent on the specific version of software installed on the SMW.

Table 1. Job Event RPM

Cray Job Event RPM
cray-jobevent_generic-<version>.rpm

RPM name varies depending on the software version. For example: cray-jobevent_generic-1.3.3-1.0502.3a5e398.3.1.ari.x86_64.rpm

Procedure

1. Copy the job event RPM located in the installation directory in `/root/sma-install/rpm/jobevent-rpm/` to the crayadm users home directory on the SMW.

```
hostname# cd /root/sma-install/rpm/jobevent-rpm/
hostname# ls cray-jobevent*
cray-jobevent_generic<version>.rpm
hostname# scp cray-jobevent_generic-<version>.rpm crayadm@smw-hostname:
```

2. Install the appropriate jobevent RPM on the SMW. The jobevent daemon will be installed in `/usr/bin/jobevent`.

Initial install:

```
smw# rpm -ihv cray-jobevent_generic-<version>.rpm
```

Upgrade:

```
smw# rpm -Uhv cray-jobevent_generic-<version>.rpm
```

Follow configuration steps according to which version of SMW software is installed.

7.1 Configuration for CLE 5.2

Prerequisites

CLE 5.2 version of SMW is installed.

The RPM has already installed the jobevent daemon as `/usr/bin/jobevent`.

Procedure

1. Copy and rename startup script to `/etc/init.d`.

```
smw# cp /opt/cray/sma/jobstats/rc.jobevent /etc/init.d/jobevent
smw# ln -s /etc/init.d/jobevent /etc/init.d/rc5.d/S15jobevent
smw# ln -s /etc/init.d/jobevent /etc/init.d/rc5.d/K01jobevent
```

2. Copy and rename configuration file to `/etc/sysconfig`.

```
smw# cp /opt/cray/sma/jobstats/sysconfig.jobevent \
/etc/sysconfig/jobevent
```

3. Define jobevent parameters in `/etc/sysconfig/jobevent` text file.

- a. Set `ENABLED` to 'yes' to enable the jobevent service.

```
ENABLED=yes
```

- b. Define the IP address of the View server (Kafka Broker).

The port number is the default Kafka port (9092).

Single Server Configuration:

```
BROKER_LIST="172.30.76.13:9092"
```

Use a space to delimitate multiple View servers.

Multiple Server Configuration:

```
BROKER_LIST="172.10.16.10:9092 172.10.16.11:9092 172.10.16.12:9092"
```

- c. Leave `TOPIC` set to 'metrics,job_events'.

```
TOPIC="metrics,job_events"
```

4. Restart daemon.

```
smw# /etc/init.d/jobevent restart
```

On an SMW-HA cluster, install rpm and configure on both nodes.



CAUTION: This procedure has not been tested for the SMW-HA cluster.

7.2 Configuration for CLE 6.0

Prerequisites

CLE 6.0 version of SMW is installed.

Procedure

1. Install the systemd unit file.

```
smw# cp /opt/cray/sma/jobstats/jobevent.service /usr/lib/systemd/system
```

2. Copy and rename the configuration file to `/etc/sysconfig`.

```
smw# cp /opt/cray/sma/jobstats/jobevent.sysconfig \
/etc/sysconfig/jobevent
```

3. Define jobevent parameters in the `/etc/sysconfig/jobevent` text file.

```
smw# vi /etc/sysconfig/jobevent
```

- a. Set `ENABLED` to 'yes' to enable the jobevent service.

```
ENABLED=yes
```

- b. Define the IP address of the View server (Kafka Broker).

The port number is the default Kafka port: 9092.

Single Server Configuration:

```
BROKER_LIST="172.10.16.10:9092"
```

Use a space to delimitate multiple View servers.

Multiple Server Configuration:

```
BROKER_LIST="172.10.16.10:9092 172.10.16.11:9092 172.10.16.12:9092"
```

- c. Leave `TOPIC` set to 'metrics,job_events.'

```
TOPIC='metrics,job_events'
```

4. Reload the systemd.

```
smw# systemctl daemon-reload
```

5. Enable and start the jobevent daemon.

```
smw# systemctl enable jobevent
smw# systemctl start jobevent
```

On an SMW-HA cluster, install rpm and configure on both nodes.



CAUTION: This procedure has not been tested for the SMW-HA cluster.

8 Add Initial ClusterStor to View for ClusterStor™

Prerequisites

View for ClusterStor™ has been installed.

Procedure

1. Go to the site specific url for View for ClusterStor™.
A security exception may appear. Click **Advanced**, and select **Proceed**.
2. Log into View for ClusterStor™.
User Name: admin
Password: admin
3. Select the gray tile with a plus sign.
A pop-up form will appear.
4. Fill out the form and select **Add a ClusterStor System**.
 - **Name:** System identifier name for the ClusterStor System
 - **IP address:** IP Address of active ClusterStor management node (n000)
 - **IP address:** IP Address of backup ClusterStor management node (n001)
 - **User:** Username for ClusterStor System read-only user, created in [ClusterStor System Configuration for Metrics](#) on page 9
 - **Password:** Password for ClusterStor System read-only user, created in [ClusterStor System Configuration for Metrics](#) on page 9

It should take no more than one minute for a new ClusterStor tile to appear. Metrics should appear within a few minutes once the ClusterStor system has been added.

5. Click **View for ClusterStor™** to refresh the system.
The browser's refresh button can also be used.

The ClusterStor tile will now display current information.

9 Configure Notification Email

Prerequisites


View for ClusterStor™ has been installed and configured.

About this task

View for ClusterStor™ creates alarms that records the state of the system as indicated by the values in specified metric series. A notification email is a message sent automatically when the state of an alarm changes. The following procedure configures

Procedure

1. Navigate to the drop-down menu in the upper right corner of the Landing Page.
2. Select **Alarm Definitions**.
3. Select the **Edit** button located at the top right of the Alarm Definitions table.



4. Enter the email address of the intended recipient of notifications.

The screenshot shows a dialog box titled 'Notification Email' with a close button (X) in the top right corner. Inside the dialog, there is a label 'Address' and a text input field containing 'root@localhost'. To the right of the input field, there is a message: 'Notifications will be sent by email to this address whenever the state of an alarm changes.' At the bottom right of the dialog, there are two buttons: 'Cancel' and 'Change Address'.

5. Click **Change Address**.

10 Change Default Password

Prerequisites

View for ClusterStor™ has been installed and a ClusterStor has been added.

About this task

For security purposes, the password should be changed from its default configuration immediately after initial installation. Changing the password is a two step process.

Procedure

CHANGE THE PASSWORD IN GRAFANA

1. Navigate to Grafana by clicking the ClusterStor title in tile.
2. Select the main menu by clicking the Grafana icon in the top left part of the screen.
3. Select **Admin** and click **Profile**.

This will take you the 'User Profile' page in Grafana.

4. Select **Change Password**.

Create a secure password.

CHANGE THE PASSWORD IN VIEW FOR CLUSTERSTOR™

5. Login into the View for ClusterStor™ as `root`.
6. Navigate to `/etc/sma-data/etc`.

```
hostname# cd /etc/sma-data/etc
```

7. Change the Grafana password in the configuration file, `site_config.yaml`.

```
hostname# vi site_config.yaml
grafana_params: {password: admin, user: admin}
```

The password has been changed.

11 Update View for ClusterStor™ Software

Prerequisites

View for ClusterStor™ has been previously installed.

Procedure

1. Stop any running containers that are not View for ClusterStor.
During the update, all Docker images will be unloaded, including any images that are not View for ClusterStor.
2. Copy software package to `/root`.
3. Unpack software package with `overwrite` option.

```
hostname# tar xvf sma-1.0.1-201805211200.tgz --overwrite
./
./sma-install/
./sma-install/sma-1.0.1-rpms-201805211200.tgz
./sma-install/setup.sh
./sma-install/templates/
./sma-install/templates/.env
./sma-install/templates/site_config.yaml
./sma-install/templates/sma.conf
./sma-install/templates/sma.service
./sma-install/templates/docker-compose.yml
```

4. Change directories to `sma-install`.

```
hostname# cd sma-install/
```

5. Run installation script with update option.

The time-date stamp in the software package file name will match the time-date stamp in the RPM package file name. For example, `sma-1.0.1-201805211200.tgz` matches with `sma-1.0.1-rpms-201805211200.tgz`.

```
hostname# ./setup.sh -u sma-1.0.1-rpms-201805211200.tgz
Updating SMA software
RPM package: sma-1.0.1-rpms-201805211200.tgz

Thu Jan  4 14:10:19 CST 2018: Stopping SMA containers...

Thu Jan  4 14:10:19 CST 2018: Checking docker and docker-compose versions...

Thu Jan  4 14:10:20 CST 2018: Prompting user for site configuration
information...
```

12 Update Retention Policies for InfluxDB

About this task

The Influx database must be modified to the duration that metrics are preserved in the database. These retention policies are highly dependent on a site's job workload and the number of OSTs the workload is distributed over. Each site will need to set the appropriate length of the InfluxDB retention policy. The default retention policy is 3 days. For the first 3 days, the data will be retained. On the fourth day, the first day's data will be pruned. On the fifth day, the second day's data will be pruned, and so on.

Retention policies can be modified at any time. For more information, see <https://docs.influxdata.com/>.

Procedure

1. Change to the SMA configuration directory.

```
hostname# cd /etc/sma-data/etc
```

2. Open the `.env` configuration file in an editor.

```
hostname# vi .env
```

3. Change the `RETENTION_POLICY_DURATION` variable to the desired value.

Note that days, months, and years are respectively defined by "d," "m," and "y."

4. Restart SMA.

```
hostname# systemctl restart sma
```