CLI "show" Commands

cvpi status <component>/all [-v=3]

-shows running, disabled, and failed components. It will show components that are failing. -v=3 adds verbosity.

cvpi resources [-v=3]

-shows memory, storage, disk throughput (>20MBps min for healthy disk, >40MBps recommended), CPUs, and NTP sync (mandatory for multi-node, at least ntpd UP for single-node)

cvpi deps <component> <start/stop>

gives the dependencies for the component to be able to start/stop

cvpi debug

-collects logs for all components for troubleshooting; collect on the primary node.

cvpi logs <component>

-to find where logs are located for a particular component. i.e. 'cvpi logs aeris' shows you /**cvpi**/**apps**/**aeris**/**logs**/. Also good for finding which node a component and subsequent logs can be found. i.e. 'cvpi logs turbine-rate-intf-counters' shows it resides on the tertiary node and its path.

cvpi info <component>

-great command to learn about a component; includes actions that can be taken, ports used, config, logging, etc.

cvpi status all -v=3 | grep disabled

-to see which processes are disabled

history

-shows list of all commands run

cvpi version

-shows version of CVP

cvpi env or cat /etc/cvpi/env

-shows environmental variables and if they are correctly set

cvpi check all

-checks that everything is set up correctly; confirms nodes are talking to each other and have same configs/env/etc.

dmesg -T

CLI "show" Commands (cont)

-shows kernel message buffer for checking disk/storage issues

CLI "Config" Commands

cvpi start/stop <component>/all

-starts/stops all available/specified components

cvpi -v=3 start/stop <component>/all

-starts/stops all available/specified components with verbosity (detail regarding failures if subcomponents fail to start

cvpi start/stop cvpi

-starts/stops cvpi stack

cvpi reset all

-resets the CVP app to its initial state via deleting all HBASE and Hadoop data

cvpi reset aeris

-deletes all Telemetry data; can be used for expedited upgrades from 2018.2.X to 2019.1.X

cvpReInstall

-case-sensitive; in the event of an install failure, execute on primary node to set all 3 nodes back to default.

cvpi config <component>/all

-configures the components

cvpi backup cvp

-new backup procedure in 2018.2.X and on

cvpi restore cvp cvp..tgz cvp.eosimages..tgz

-new restore procedure in 2018.2.X and on; can't restore across major releases due to data formatting changes (i.e. can't restore from 2018.X to 2019.X)

cvpi enable cvpi

-enables components of CVP to be automatically restarted if they stop

cvpi init

-gets rid of corrupted data folders; recreates directory structure; repairs any damage by removing whole directories

hdfs dfsadmin -safemode get

-checks to see if hadoop/hbase in safe mode

CLI "Config" Commands (cont)

hdfs dfsadmin -safemode leave

*-try to get primary/secondary to leave **safe** mode; then try to start it again*

hdfs hbck

-checks for inconsistencies/corruptions; prints OK or gives Errors; run several times as some inconsistencies are transient

hdfs hbck -repair

-repair inconsistencies; run 5-10 times if necessary

/cvpi/zookeeper/bin/zkServer.sh start/stop

-if seeing zookeeper issues; zookeeper won't be stopped via 'cvpi stop all'

systemctl stop cvpi-watchdog.timer

In a cluster, will need to stop the watchdog timer when stopping zookeeper on all three nodes otherwise it will spawn a new zookeeper process.

MINIMUM Requirements

Lab (<25 devices)	Production (<=500 devices)
CPUs: 16 cores	CPUs: 16 cores
RAM: 16GB	RAM: 22GB
Disk: 125GB	Disk: 1 TB
Disk Throughput: 20MB/s	Disk Throughput: 40++MB/s

More might be needed based on feature sets in use. For example:

For CloudVision Wifi:
+4 CPU
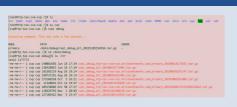+8 GB RAM
+100GB Disk storage
+10 charisma

For Elasticsearch (MAC/IP search feature):
+4 CPU

Also for Production, 16 Cores could be 8 CPU x 2 Core or 16 CPU x1 Core.

Where are the debug files?

Device/Interface Scale (multi-node cluster)

	Supported Scale Target	Datacenter (80 intfs/device)	OR Campus POE (50 intfs/device)	OR vEOS (8 intfs/device)
Foster/2018.2.3 (SHIPPED)	500 devices + 20K intfs	250	400	500
Grant/2019.1.0 (Sept, EFT)	500 devices + 20K intfs `Updated`	250	400	500
2019.1.x (Q4, TESTING)	1000 devices + 40K intfs	500	800	1000

As customers close in on these numbers, expect give and take with additional beta features, latency, etc. as resources reach capacity.

Where is it?

From root ==> su cvp ==> /cvpi

all scripts, packages, config files, logs

Logs

/cvpi/logs; /cvpi/hbase/logs; /cvpi/hadoop/logs; /cvpi/tomcat/logs

Shortcut to logs

Also just run $ cvpi logs <component> which shows path to logs.

Config Files

/cvpi/conf/components/; /cvpi/apps/turbine/configs/; /cvpi/apps/aeris/conf/; /cvpi/apps/cvp/conf/; /cvpi/apps/geiger/conf/; /cvpi/apps/wifimanager/conf

Backups

*/**data**/**cvpbackup**/ on the **primary**; backups are run nightly at 2am UTC by default; check via **crontab -l** as root user; 5 backups stored*

Minimum Configuration on EOS Device

```
Confirm the daemon is correctly installed.
!
daemon TerminAttr
   exec /usr/bin/TerminAttr -ingestgrpcurl=10.8-
1.110.104:9910 -cvcompression=gzip -ingestauth=-
key,cvp -smashexcludes=ale,flexCounter,hardw-
are,kni,pulse,strata -ingestexclude=/Sysdb/cell/-
1/agent,/Sysdb/cell/2/agent -ingestvrf=default -
taillogs
   no shutdown
!
```

ntpd needs to be enabled for single node; NTP sync essential for multi-node.

```
!
ntp server 10.81.111.240 prefer iburst
ntp server 10.81.111.241 iburst
!
```

By **sh-arista**
cheatography.com/sh-arista/

Not published yet.
Last updated 22nd October, 2020.
Page 2 of 3.

Sponsored by **ApolloPad.com**
Everyone has a novel in them. Finish Yours!
https://apollopad.com

Minimum Configuration on EOS Device (cont)

```
Turn up api for http for EAPI to work; turn up
unix-socket so TerminAttr can talk to ConfigAgent
(nginx method).
management api http-commands
    protocol http
    protocol unix-socket
    no shutdown
!
```

TerminAttr has 2 mechanisms to talk to ConfigAgent:
Default VRF - via unix socket directly, no additional config required
Non-default VRF - cannot talk directly (ConfigAgent only listens in the Default VRF) so the connection has to go via nginx; protocol unix-socket required under management api http-commands.

Enabling LANZ on EOS CLI

```
queue-monitor length
!
queue-monitor streaming ⇒ TerminAttr runs in
default VRF so this has to be in default as well!
no shutdown
!
```

Can confirm in bash via **curl localhost:6060/rest/LANZ/congestion**