RHEL6: Native Defense In Depth Tooling

Shawn Wells
Director, Innovation Programs
shawn@redhat.com

Norman Mark St. Laurent,
Senior Solutions Architect
msl@redhat.com

Agenda

- Security Compliance Management [Prevention]
  Shawn Wells

- Incident Response [Detection + Response]
  Mark St Laurent

SCAP Security Guide

- Delivers practical security guidance, baselines, and associated validation mechanisms using the SCAP protocol suite.

- Current upstream source for STIG and NSA SNAC documents

SCAP Security Guide

- Recommendations map to institutionalized policies where applicable
- Because of this mapping, we can create custom profiles
  - RHEL6 STIG (collaboration with DISA FSO)
  - RHEL6 SNAC (collaboration with NSA)
  - Baseline content for NIST 800-53 (working on USGCB now!)

Sample SSG Rule

```
<Rule id="enable_auditd_service”>
    <title>Enable auditd Service</title>
    <description>
      The <tt>auditd</tt> service is an essential userspace
      component of the Linux Auditing System, as it is
      responsible for writing audit records to disk.
      <service-enable-macro service="auditd" />
    </description>
    <ocil><service-enable-check-macro service="auditd" /></ocil>
    <rationale>
      Ensuring that the <tt>auditd</tt> service is active ensures that
      audit records generated by the kernel can be written to disk, or
      that appropriate actions will be taken if other obstacles exist.
    </rationale>
    <ident cce="4292-9" />
    <oval id="service_auditd_enabled" />
    <ref nist="CM-6, CM-7"
disa="169,172,174,1353,1462,1487,1115,1454,067,158,831,1123,1190,1312,126
3,130" />
</Rule>
```

Sample SSG Check

```
<ind:textfilecontent54_object id="obj_20134" version="1">
    <ind:path>/etc</ind:path>
    <ind:filename>sysctl.conf</ind:filename>
    <ind:pattern operation="pattern match">^\s*net\.ipv6\.conf\.all
\.disable_ipv6\s*=\s*1$</ind:pattern>
    <ind:instance datatype="int">1</ind:instance>
 </ind:textfilecontent54_object>
```

Sample Scan Results

Title	Result
Ensure /tmp Located On Separate Partition	fail
Ensure /var Located On Separate Partition	fail
Ensure /var/log Located On Separate Partition	fail
Ensure /var/log/audit Located On Separate Partition	fail
Ensure /home Located On Separate Partition	fail
Encrypt Partitions	notchecked
Ensure Red Hat GPG Key Installed	pass
Ensure gpgcheck Enabled In Main Yum Configuration	pass

Sample Scan Results

Result for Ensure gpgcheck Enabled In Main Yum Configuration

Result: **pass**

Rule ID: **ensure_gpgcheck_globally_activated**

Time: **2013-03-23 13:43**

Severity: **high**

The `gpgcheck` option should be used to ensure checking of an RPM package's signature always occurs prior to its installation. To configure yum to check package signatures before installing them, ensure the following line appears in `/etc/yum.conf` in the `[main]` section:

```
gpgcheck=1
```

Ensuring the validity of packages' cryptographic signatures prior to installation ensures the provenance of the software and protects against malicious tampering.

Security identifiers

- CCE-26709-6

results overview

RHN Satellite Integration

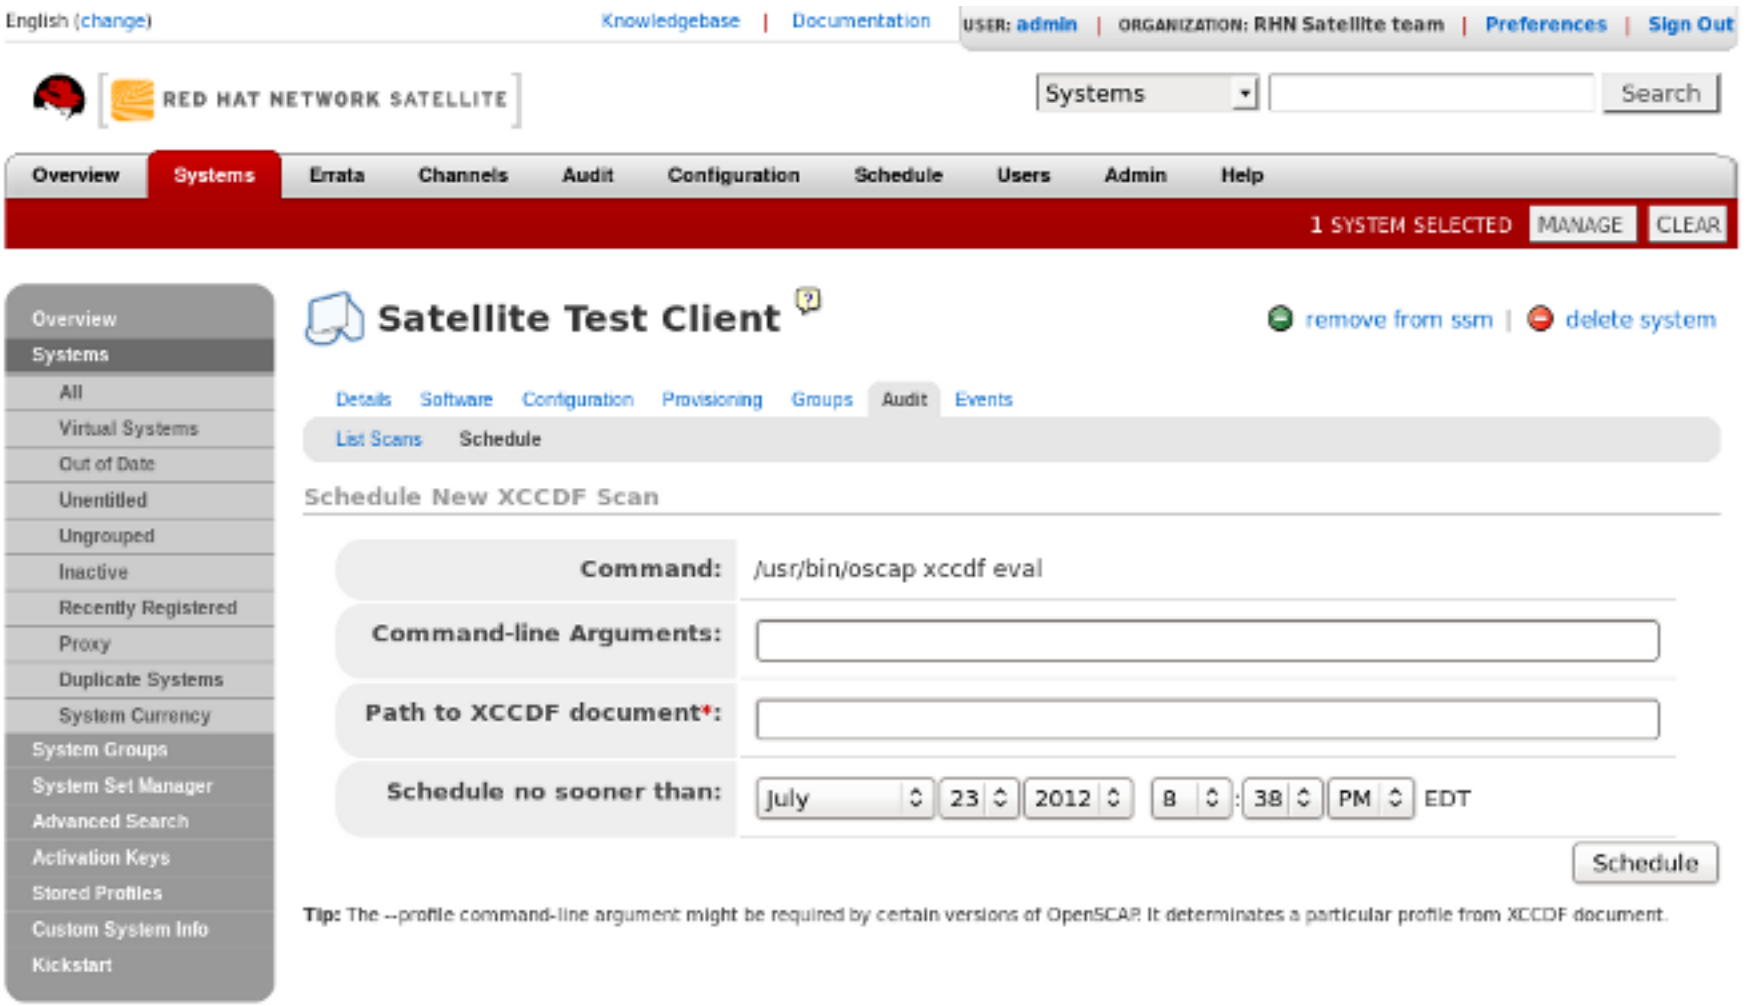

RHN Satellite Integration

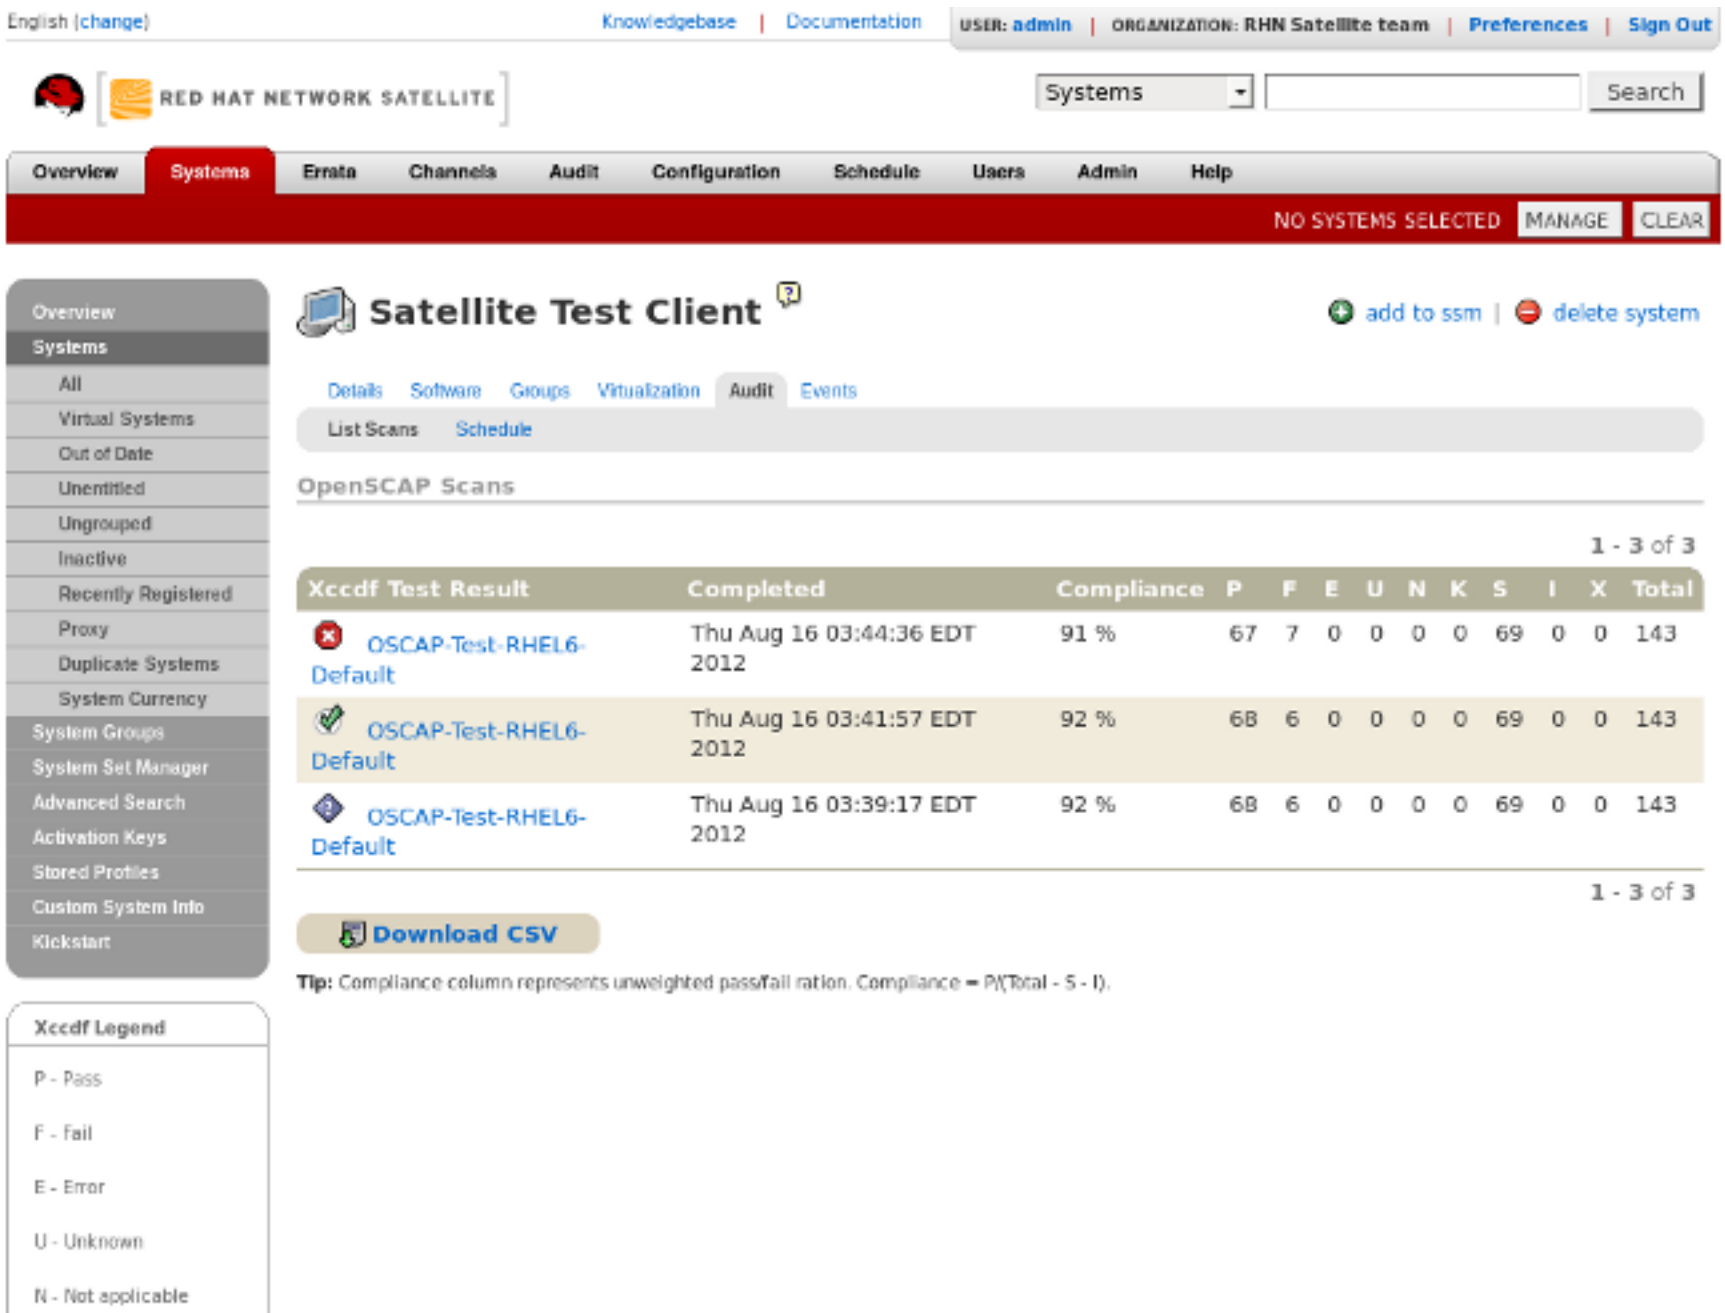

More Information

- Project Homepage
  https://fedorahosted.org/scap-security-guide/

- Mailing List
  https://fedorahosted.org/mailman/listinfo/scap-security-guide

redhat.

USING AN OPEN SOURCE FRAMEWORK TO CATCH THE BAD GUY

BUILT-IN FORENSICS, INCIDENT RESPONSE, AND SECURITY WITH RED HAT ENTERPRISE LINUX 6

Norman Mark St. Laurent,
Senior Solutions Architect
Red Hat, Inc.
msl@redhat.com

Agenda

- Integrity Checking with RHEL 6
  - Background
  - Terminology
- Operational Use of RHEL 6 RPM and AIDE
  - RPM
  - AIDE
- Conclusion

Agenda

- **Integrity Checking with RHEL 6**
  - Background
  - Terminology
- Operational Use of RHEL 6 RPM and AIDE
  - RPM
  - AIDE
- Conclusion

Agenda

- **Integrity Checking with RHEL 6**
  - **Background**
  - Terminology
- Operational Use of RHEL 6 RPM and AIDE
  - RPM
  - AIDE
- Conclusion

Integrity Checking with RHEL 6

- **Part 1:**

  Using the RHEL audit sub system for forensics and incident response to meet security requirement objectives and goals.

  Closely followed and mapped NIST Special Publication 800-92 “Guide to Computer Security Log Management”.

  http://www.redhat.com/solutions/industry/government/

RED HAT ENTERPRISE LINUX 6

Integrity Checking with RHEL 6

- **Part 2:**

  Integrity checking with RHEL 6, involves calculating a message digest for each file and storing the message digest securely to ensure changes to files such as archived logs are detected.

  A message digest, also called a digital signature, uniquely identifies data and has the property that changing a single bit in the data causes a completely different message digest to be generated.

http://www.redhat.com/solutions/industry/government/

Integrity Checking with RHEL 6

- RHEL 6 allows incident response, forensics examiners, and system administrators easy access to lightweight, easy-to-use tools and techniques that allow them to quickly identify file system modifications, changes or compromises.

- A host-based IDS provides the data integrity needed to ensure adequate protection of information and system data, and helps meet security requirements and compliance.

Integrity Checking with RHEL 6

- RPM Package Manager (RPM) and Advanced Intrusion Detection Environment (AIDE) delivers continuous and automated monitoring for security compliance and for implementing the needed security controls for a true “Defense in Depth” approach.
  - Revealing unwanted changes.

Integrity Checking with RHEL 6

- An altered environment could be achieved on a compromised system where a ***root kit*** has been installed.
  - A root kit contains a collection of “Bad Guy” tools that an intruder installs on a victims computer after gaining access.
    - Log-cleaning scripts, network sniffers, and most importantly Trojan replacements of core system utility programs....

      last, ps, netstat, killall, ifconfig, find, du, passwd, pidof, tcpd, top...

Integrity Checking with RHEL 6

- A root kit does not have to be a sophisticated program. It could be as easy as replacing the binary of the last program with a simple wrapper script.

#last | awk ‘$1 !~ /UserName/ { print 0}’

Integrity Checking with RHEL 6

- RPM can be used as a host-based IDS. Verification options with RPM can be invaluable to a forensics investigation and could detect critical system files and executables that have been modified.

- AIDE is a file integrity checker. Using rules read from /etc/aide.conf AIDE will create a database of file attributes and extended attribute information. Once the database is initialized, it can be used to verify the integrity of files.
  - Hashing Algorithms: md5, sha1, rmd10, tiger, haval, sha256, and sha512.

Integrity Checking with RHEL 6

- Every Good Computer Forensics Examiner / Incident Response Analyst understands **atime**, **ctime**, and **mtime**.....

Agenda

- **Integrity Checking with RHEL 6**
  - Background
  - **Terminology**
- Operational Use of RHEL 6 RPM and AIDE
  - RPM
  - AIDE
- Conclusion

Integrity Checking with RHEL 6

- It is important to understand specific terms dealing with Intrusion detection tools with RHEL 6.
- Understanding the terms will help you if your system defenses has been penetrated or if you come in after-the-fact for forensics analysis.
  - They will allow you to know the factors and mechanics needed to:
    - Perform a forensics exam
    - Deep dive into an incident
    - Assist when something goes wrong from a systems administrative perspective.

Integrity Checking with RHEL 6

- **Baseline**

A baseline is a set of data used for comparison. A major advantage of a baseline is that a file can be summarized by its file attributes.

When file attributes provide enough detail about a fiven file, then changes to that file can be detected by comparing snapshot data to the corresponding baseline data.

A baseline is usually the first snapshot taken.

Integrity Checking with RHEL 6

- **Snapshot**

A snapshot is a view of something transitory in nature. A snapshot can be taken multiple times to provide a “picture” of a specific point in time.

In the context of file integrity, a snapshot refers to a set of critical observations (file attributes) taken at a particular point in time from one or more file objects stored on a given Red Hat system.

When a particular snapshot is used as the reference in cange analysis, it is called a baseline.

Integrity Checking with RHEL 6

- **Change Analysis**

Change analysis is the process of comparing a snapshot to a baseline.

Change analysis improves the overall security and operational performance by automatically analyzing detected changes against security policies.

It enhances file integrity monitoring to close the time gap between a bad change of a file and detecting and correcting it.

Integrity Checking with RHEL 6

- **Digital Signature | Message Digest | Hash**

A digital signature (also known as message digest, cryptographic check-sum or cryptographic hash) is nothing more than a number.

A special number that is effectively a hash code produced by a functin that is very difficult to reverse.

A message digest is a cryptographic hash computed over a given block of data such as a file's content.

Integrity Checking with RHEL 6

- **Digital Signature | Message Digest**

Integrity: A digital signature indicates whether a file or a message has been modified.

Authentication: A digital signature allows you to mathematical verify the name of the file.

Agenda

- Integrity Checking with RHEL 6
  - Background
  - Terminology
- **Operational Use of RHEL 6 RPM and AIDE**
  - RPM
  - AIDE
- Conclusion

Agenda

- Integrity Checking with RHEL 6
  - Background
  - Terminology
- **Operational Use of RHEL 6 RPM and AIDE**
  - **RPM**
  - AIDE
- Conclusion

Integrity Checking with RHEL 6

- Engineers at Red Hat and in the open source community developed RPM Package Manger.
  - Used to Install ,upgrade, and verify software packages on a Red hat system.
  - The verify feature of RPM can be used to check file integrity and make sure that files have not been modified or replaced.
  - Verifying a package compares information about the installed files in the package with information about the files taken from the package metadata stored in the RPM database.

Integrity Checking with RHEL 6

- The RPM database is known as the Baseline.
- File size, file type, owner and group, file permissions, time stamps and the MD5 are all stored in the database.
- The format of the rpm -V or rpm --verify command is a string of characters. The mnemonically em**B**oldened character denotes failure of the corresponding test.

  A single “.” (period) means the test passed, while a single “?” means the test could not be performed (e.g., file permissions prevent reading).

Integrity Checking with RHEL 6

Output String	Definition
S	The file **S**ize differs
M	**M**ode differs (includes permissions and file type)
5	MD**5** sum differs
D	**D**evice major/minor number mismatch
L	read**L**ink (2) mismatch
U	**U**ser ownership differs
G	**G**roup ownership differs
T	m**T**ime differs
P	ca**P**abilities differ

Integrity Checking with RHEL 6

- In the root kit example, lets say you need to find out which Red Hat package the pidof command is a part of.

  #rpm -qf `which pidof`
  sysvinit-tools-2.87.4.dsf.el6.x86_64

  #rpm --verify sysvinit-tools-2.87.4.dsf.el6.x86_64

Integrity Checking with RHEL 6

-

```
[root@mstlaure mstlaure]# rpm -qf /etc/profile
setup-2.8.14-20.el6.noarch
[root@mstlaure mstlaure]# rpm --verify setup
SM5....T.  c /etc/hosts.allow
S.5....T.  c /etc/printcap
S.5....T.  c /etc/profile
[root@mstlaure mstlaure]#
```

Integrity Checking with RHEL 6

- An additional RPM challenge is that an intelligent attacker will modify RPM itself to hide changes done by a root kit.
  - Allowing the attacker to mask its intrusion and gain root privilege.

    Note: To solve this, you should implement a secondary check that can also be run completely independent of the installed system.

    ***RPM from a Boot Disk as an Example...

  - RPM was not designed to specifically be a file auditing tool. However, the nice thing is that it is installed on every version of supported RHEL.

Integrity Checking with RHEL 6

- In rescue mode, you can use RPM to install, remove, update or verify packages from the installed system:

1. Boot the system into rescue mode with the linux rescue command at the boot prompt.
2. Change to the root directory

   #chroot /mnt/sysimage

3. Use RPM to verify a package

   #rpm --verify /package/of/your/choice.rpm

Integrity Checking with RHEL 6

- Backing up the RPM database

  The **/var/lib/rpm** directory stores all files used by the RPM command. All files inside this directory are binary files.

[root@mstlaure mstlaure]# file /var/lib/rpm/*
/var/lib/rpm/Basenames: Berkeley DB (Hash, version 9, native byte-order)
/var/lib/rpm/Conflictname: Berkeley DB (Hash, version 9, native byte-order)
/var/lib/rpm/__db.001: Applesoft BASIC program data
/var/lib/rpm/__db.002: 386 pure executable
/var/lib/rpm/__db.003: 386 pure executable
/var/lib/rpm/__db.004: 386 pure executable
/var/lib/rpm/Dirnames: Berkeley DB (Btree, version 9, native byte-order)
/var/lib/rpm/Filedigests: Berkeley DB (Hash, version 9, native byte-order)

Integrity Checking with RHEL 6

- To back up the RPM data base stored in the /var/lib/rpm directory:

  1. Log in as the root user.
  2. Remove any Stale Lock Files. (You may need to make ensure that no application has any of the RPM database files open using the **lsof** command).

  #losf /var/lib/rpm
  #/bin/rm -f /var/lib/rpm/__db*

  3. Backup the /var/lib/rpm directory using the **tar** command. Once the back up is complete move the tar file to safe offline storage.

  #tar -czfg $ (hostname).rpmdatabase.tar.gz /var/lib/rpm

Integrity Checking with RHEL 6

- A couple of neat RPM scripts will be presented at this years Red Hat Summit in Boston.
  - A script that will find all files in a given directory (or on a system from “/” that do not belong to an RPM.

    Note: It is important to verify which files were not put on a system by an RPM file.
  - A script that can be used as a System Integrity checking script (finding all files on a system that have been modified).

Agenda

- Integrity Checking with RHEL 6
  - Background
  - Terminology
- **Operational Use of RHEL 6 RPM and AIDE**
  - RPM
  - **AIDE**
- Conclusion

Integrity Checking with RHEL 6

- Aide is a tool for monitoring file system changes.
- Aide was written to be a simple and free alternative to Tripwire.
- Aide is not installed by default on a RHEL 6 system.
- Aide will build a database of the files specified in the configuration file /etc/aide.conf

Integrity Checking with RHEL 6

- Aide stores various file attributes including: permissions, inode number, user, group, file size, mtime-ctime-atime, growing size, number of links, and link name.

- Aide also generates a cryptographic check-sum of each file using one or a combination of the following message digest algorithms:

  sha1, sha256, sha512, md5, rmd160, and tiger.

Integrity Checking with RHEL 6

- After installation and setting up the RHEL 6 AIDE RPM, the extended attributes acl, xattr, and selinux are ready to be used because they have been explicitly enabled during compile time.

- The configuration file is /etc/aide.conf. Setting up this file can be difficult and must be done on a case-by-case basis.

  Note: At this years Red Hat summit, we will dive deeply into the /etc/aide.conf file.

Integrity Checking with RHEL 6

- The file /etc/aide.conf contains the runtime configuration AIDE uses to initialize or check the AIDE database.

- There are 3 types of of lines in the /etc/aide.conf file:
  - Configuration lines used to set the configuration parameters and define/undefine variables.
  - Selection lines used to indicate which files are added to the database.
  - Macro lines used to define (or undefine) variables within the configuration file.

Note: Lines starting with **#** are comments.

Integrity Checking with RHEL 6

- The standard RHEL 6 /etc/aide.conf file is set up to pass the Common Criteria Labeled Security Protection Profile (LSPP).

Note: This is a huge benefit for installing on government systems and commercial systems who need to pass specific security accreditation and allow for appropriate security policy settings.

(DoD 5300, DCID 6/3 at PL3, SOX, HIPPA, and PCI)

Integrity Checking with RHEL 6

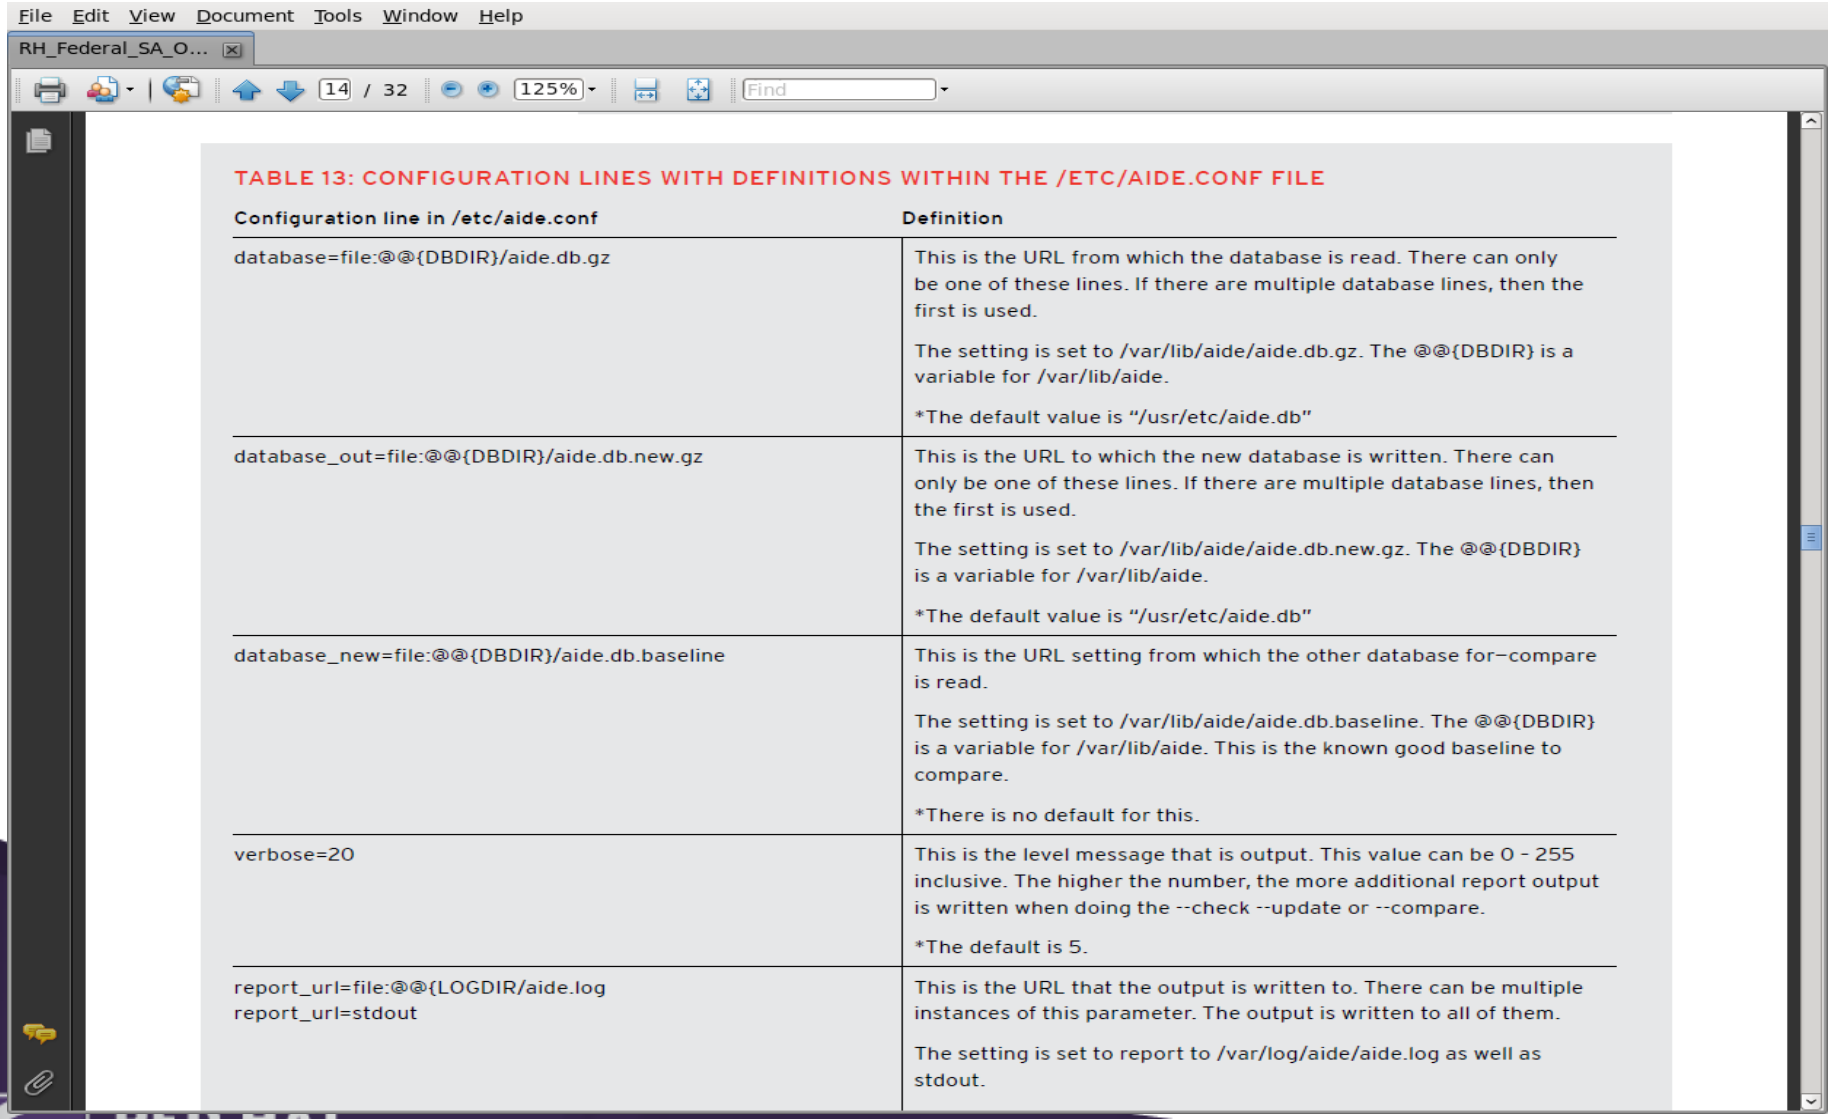

File Edit View Document Tools Window Help

RH_Federal_SA_O...

14 / 32 125% Find

TABLE 13: CONFIGURATION LINES WITH DEFINITIONS WITHIN THE /ETC/AIDE.CONF FILE

Configuration line in /etc/aide.conf	Definition
database=file:@@{DBDIR}/aide.db.gz	This is the URL from which the database is read. There can only be one of these lines. If there are multiple database lines, then the first is used. The setting is set to /var/lib/aide/aide.db.gz. The @@{DBDIR} is a variable for /var/lib/aide. *The default value is "/usr/etc/aide.db"
database_out=file:@@{DBDIR}/aide.db.new.gz	This is the URL to which the new database is written. There can only be one of these lines. If there are multiple database lines, then the first is used. The setting is set to /var/lib/aide/aide.db.new.gz. The @@{DBDIR} is a variable for /var/lib/aide. *The default value is "/usr/etc/aide.db"
database_new=file:@@{DBDIR}/aide.db.baseline	This is the URL setting from which the other database for–compare is read. The setting is set to /var/lib/aide/aide.db.baseline. The @@{DBDIR} is a variable for /var/lib/aide. This is the known good baseline to compare. *There is no default for this.
verbose=20	This is the level message that is output. This value can be 0 - 255 inclusive. The higher the number, the more additional report output is written when doing the --check --update or --compare. *The default is 5.
report_url=file:@@{LOGDIR/aide.log report_url=stdout	This is the URL that the output is written to. There can be multiple instances of this parameter. The output is written to all of them. The setting is set to report to /var/log/aide/aide.log as well as stdout.

RED HAT ENTERPRISE LINUX 6

Integrity Checking with RHEL 6

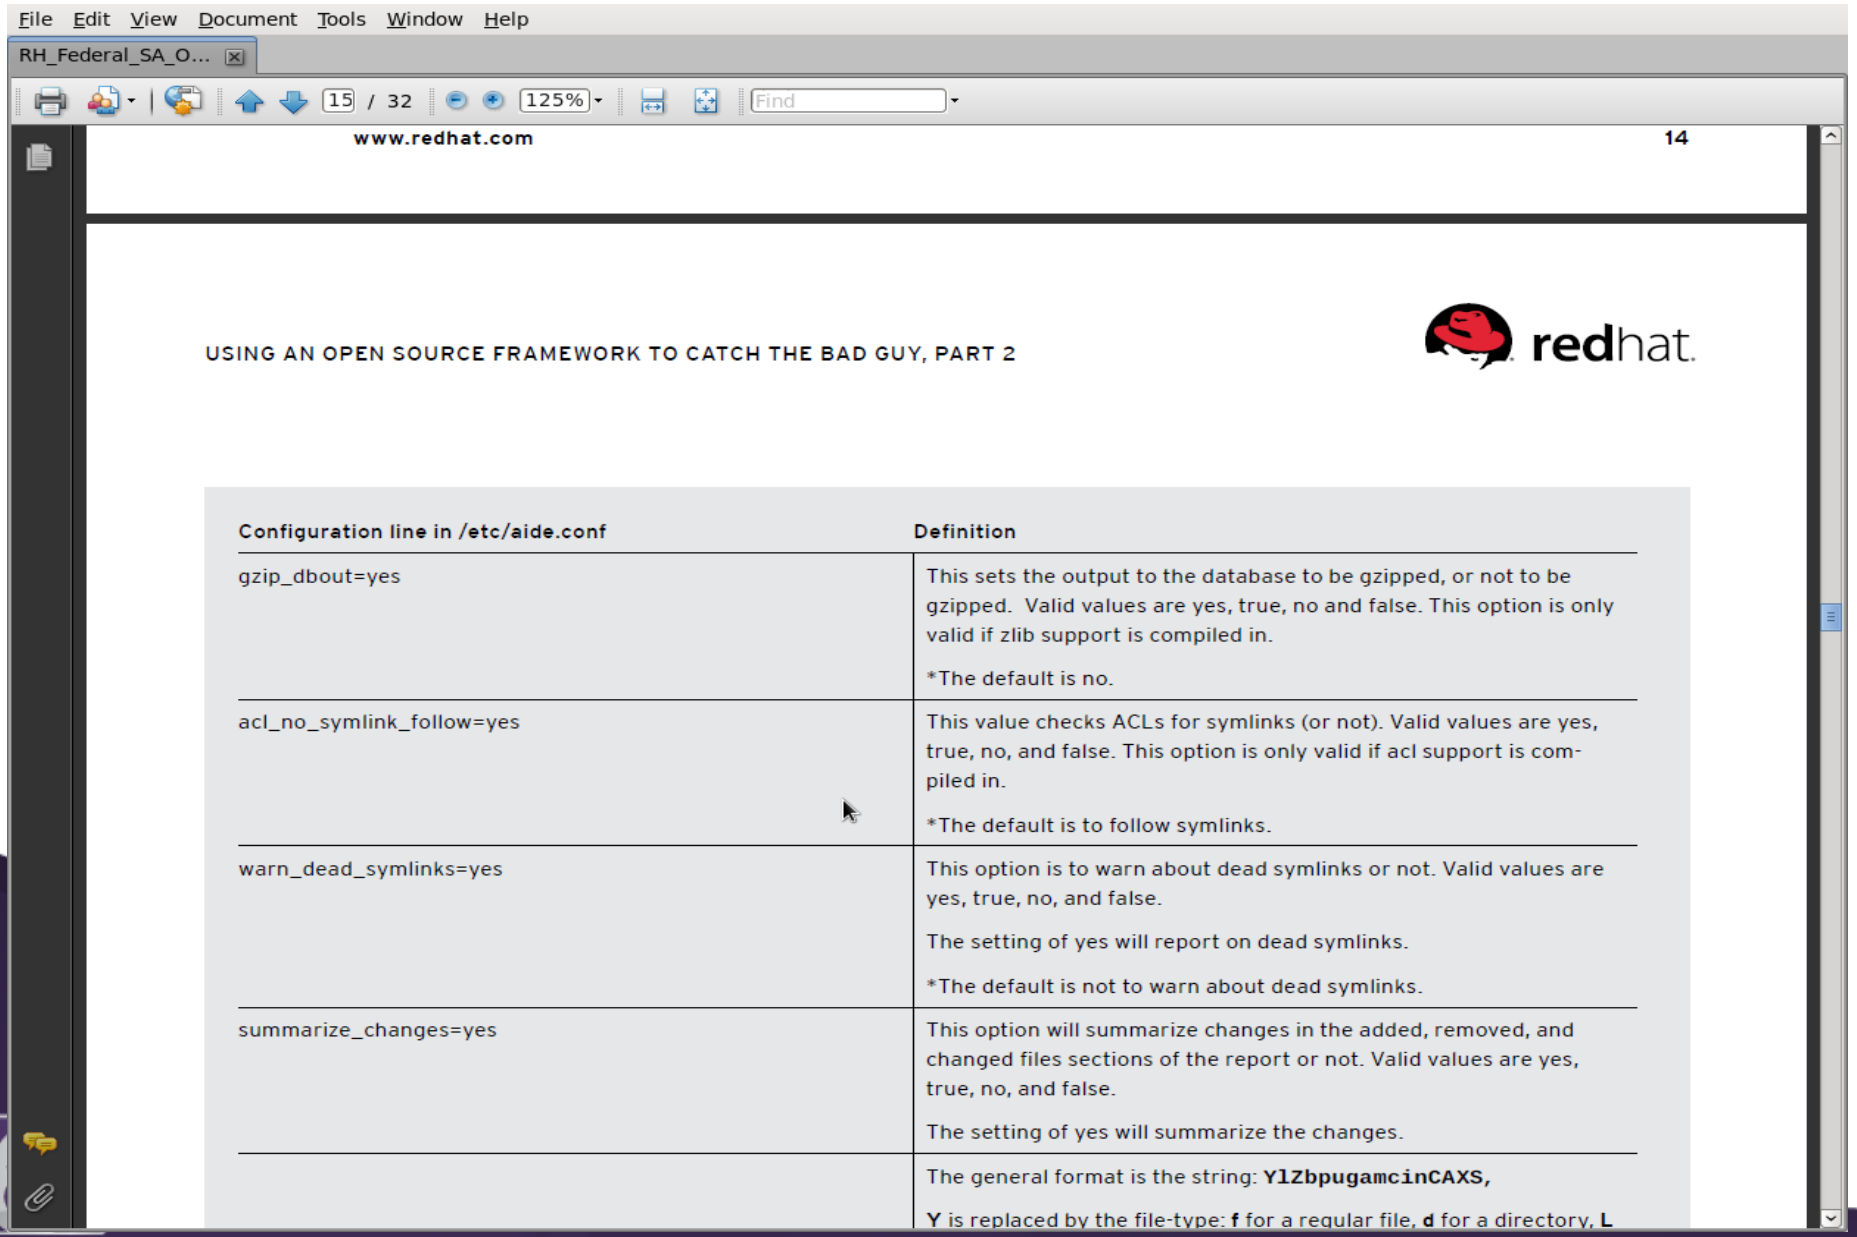

File Edit View Document Tools Window Help

RH_Federal_SA_O...

15 / 32 125% Find

USING AN OPEN SOURCE FRAMEWORK TO CATCH THE BAD GUY, PART 2

redhat.

Configuration line in /etc/aide.conf	Definition
gzip_dbout=yes	This sets the output to the database to be gzipped, or not to be gzipped. Valid values are yes, true, no and false. This option is only valid if zlib support is compiled in. *The default is no.
acl_no_symlink_follow=yes	This value checks ACLs for symlinks (or not). Valid values are yes, true, no, and false. This option is only valid if acl support is compiled in. *The default is to follow symlinks.
warn_dead_symlinks=yes	This option is to warn about dead symlinks or not. Valid values are yes, true, no, and false. The setting of yes will report on dead symlinks. *The default is not to warn about dead symlinks.
summarize_changes=yes	This option will summarize changes in the added, removed, and changed files sections of the report or not. Valid values are yes, true, no, and false. The setting of yes will summarize the changes.
	The general format is the string: **YlZbpugamcinCAXS,** **Y** is replaced by the file-type: **f** for a regular file, **d** for a directory, **L**

Integrity Checking with RHEL 6

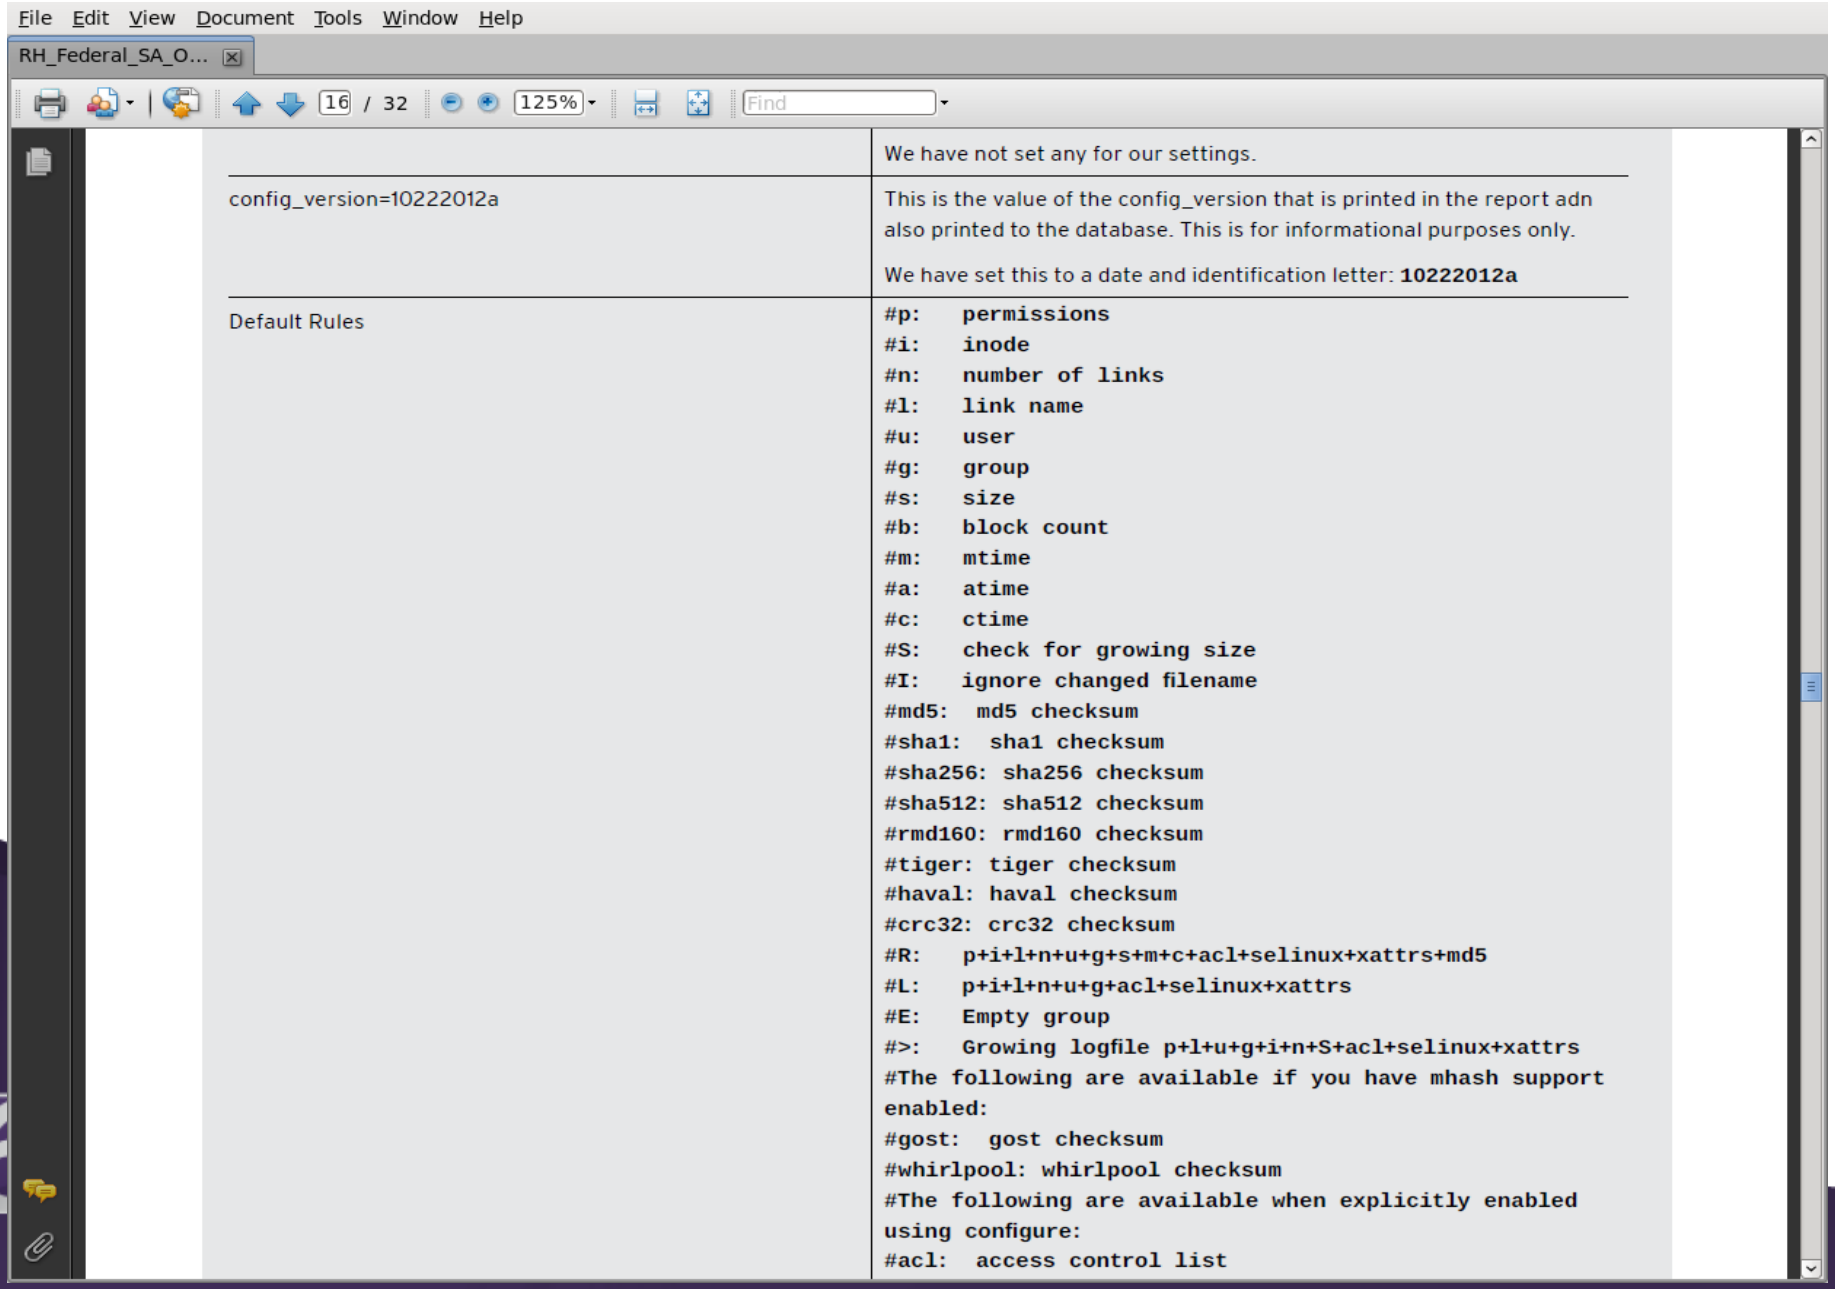

Integrity Checking with RHEL 6

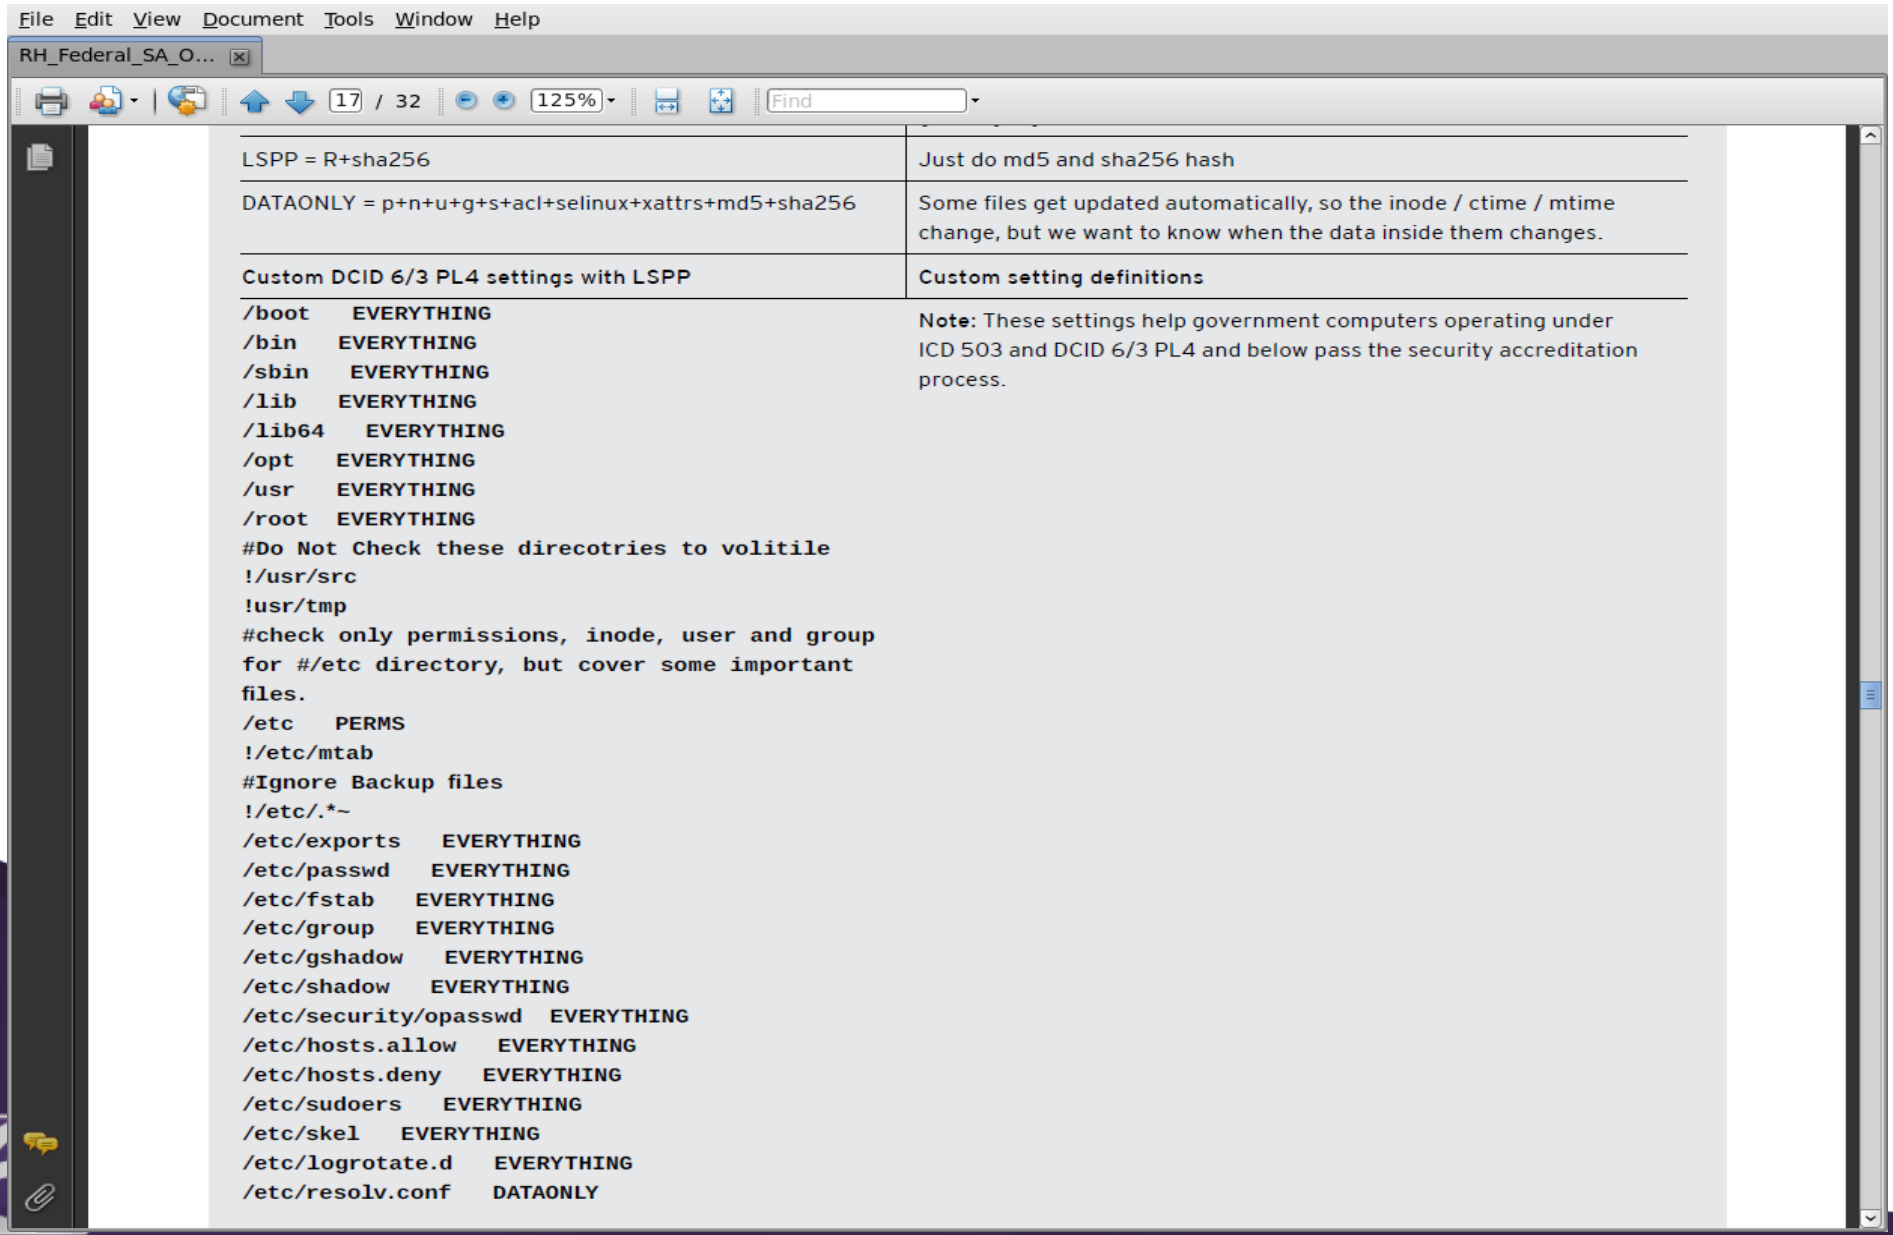
File Edit View Document Tools Window Help

RH_Federal_SA_O...

17 / 32 125% Find

LSPP = R+sha256	Just do md5 and sha256 hash
DATAONLY = p+n+u+g+s+acl+selinux+xattrs+md5+sha256	Some files get updated automatically, so the inode / ctime / mtime change, but we want to know when the data inside them changes.

Custom DCID 6/3 PL4 settings with LSPP	Custom setting definitions

```
/boot    EVERYTHING
/bin    EVERYTHING
/sbin    EVERYTHING
/lib    EVERYTHING
/lib64    EVERYTHING
/opt    EVERYTHING
/usr    EVERYTHING
/root   EVERYTHING
#Do Not Check these direcotries to volitile
!/usr/src
!usr/tmp
#check only permissions, inode, user and group
for #/etc directory, but cover some important
files.
/etc    PERMS
!/etc/mtab
#Ignore Backup files
!/etc/.*~
/etc/exports    EVERYTHING
/etc/passwd    EVERYTHING
/etc/fstab    EVERYTHING
/etc/group    EVERYTHING
/etc/gshadow    EVERYTHING
/etc/shadow    EVERYTHING
/etc/security/opasswd   EVERYTHING
/etc/hosts.allow    EVERYTHING
/etc/hosts.deny    EVERYTHING
/etc/sudoers    EVERYTHING
/etc/skel    EVERYTHING
/etc/logrotate.d    EVERYTHING
/etc/resolv.conf    DATAONLY
```

Note: These settings help government computers operating under ICD 503 and DCID 6/3 PL4 and below pass the security accreditation process.

Integrity Checking with RHEL 6

- Creating the AIDE Baseline Database
  - Once the /etc/aide.conf file has been constructed, it is time to create the baseline database.

    Note: This should be done as soon as all the system tools and programs have been loaded and installed and before it is put onto the network.

  - After specific changes have been made to the /etc/aide.conf file, you can validate that the file is capable of running (e.g., no syntax errors).

Integrity Checking with RHEL 6

- The error message contains the last successful read line of the /etc/aide.conf file. In the example here, line 40 has an unexpected error.

```
#aide --verbose=255
#aide --config-check
39:sysntax error:!
39:error while reading configuration:!
Configuration error
```

Integrity Checking with RHEL 6

- You can specify where the Baseline Database goes in the /etc/aide.conf file. Default is:

  /var/lib/aide/aide.db.gz

- As soon as the baseline database has been created, a copy should be saved to a secure location (e.g., CD-R or DVD-R or a remote server or USB Disk) for later snapshot comparisons.

Integrity Checking with RHEL 6

- Initialize the AIDE database:

  #aide -i
  #mv /var/lib/aide/aide.db.gz /mnt/usbdrive/aide_baseline_database/aide.db.gz

- Create the AIDE snapshot for comparison:

  #cp /mnt/usbdrive/aide_baseline_database/aide.db.gz /var/lib/aide/aide.db.baseline.gz

  Note: The first step in creating the snapshot comparison is to copy the saved baseline database to the file that you noted in the /etc/aide.conf. This is denoted with the

  database_new=file:@@(DBDIR)/aide.db.baseline directive.

Integrity Checking with RHEL 6

- Perform the snapshot comparison.

  #aide --check

  Note: If no output follows, then, no changes were detected.

Integrity Checking with RHEL 6

- Perform the snapshot comparison.

#aide --check

Note: If output follows, then, changes were detected.

AIDE found differences between database and filesystem!!

Summary:
Total number of files: 32,056
Added files: 0
Removed files: 0
Changed files: 1
....

Integrity Checking with RHEL 6

- Snapshot Comaprison in Verbose Mode will go into great detail about any and all changes.

#aide –check -V

AIDE found differences between database and filesystem!!

Summary:
Total number of files: 32,056
Added files: 0
Removed files: 0
Changed files: 1

Changed Files:
---
changed: /etc/passwd

Agenda

- Integrity Checking with RHEL 6
  - Background
  - Terminology
- Operational Use of RHEL 6 RPM and AIDE
  - RPM
  - AIDE
- **Conclusion**

Integrity Checking with RHEL 6

- From an intrusion detection point-of-view, incident response teams, forensic analysts, systems administrators, and security engineers need the ability to detect change correctly and consistently on their machines.

- Using lightweight, open source tools that come with RHEL 6 such as RPM and AIDE provide incident response techniques that are recognized in the court of law.

Integrity Checking with RHEL 6

- Klayton Monroe Quote:

  “When it comes to incident response and forensics analysis, one snapshot is better than none. Two snapshots are better than one, and a complete history of snapshots is nearly ideal.

  However, even if no baseline exists, one should not rule out the usefulness of baselining tools. A significant amount of information can be collected or derived from data collected by such tools even though a history of prior snapshots does not exist.”