Big problems with big data – Hadoop interfaces security

Jakub Kaluzny

OWASP 中国

The Open Web Application Security Project

whoami

OWASP 中国
The Open Web Application Security Project

Sr. IT Security Consultant at SecuRing

- Consulting all phases of development
- penetration tests
- high-risk applications and systems

Researcher

- Hadoop, FOREX, MFP printers, proprietary network protocols

Agenda

OWASP 中国
The Open Web Application Security Project

- Big data nonsenses
- Crash course on hacking Hadoop installations
- Ways to protect big data environments

Expect some CVEs

Results summary

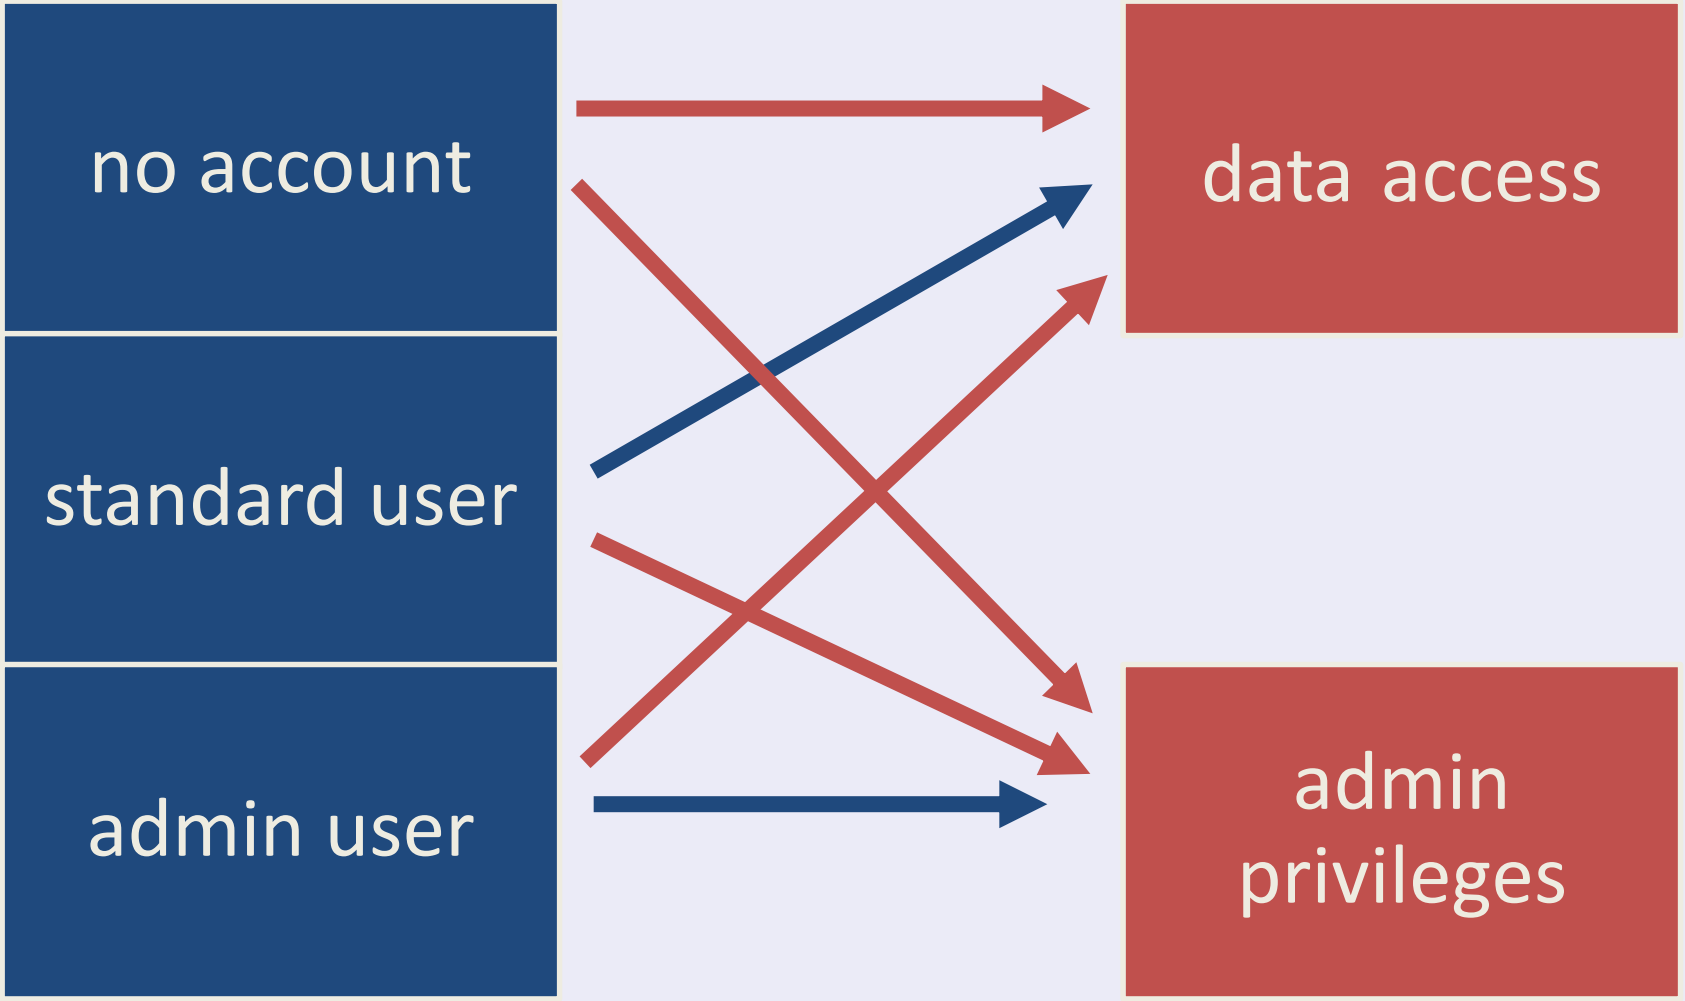

OWASP 中国
The Open Web Application Security Project

Know your target

WHAT IS HADOOP?

Normal database

OWASP 中国
The Open Web Application Security Project

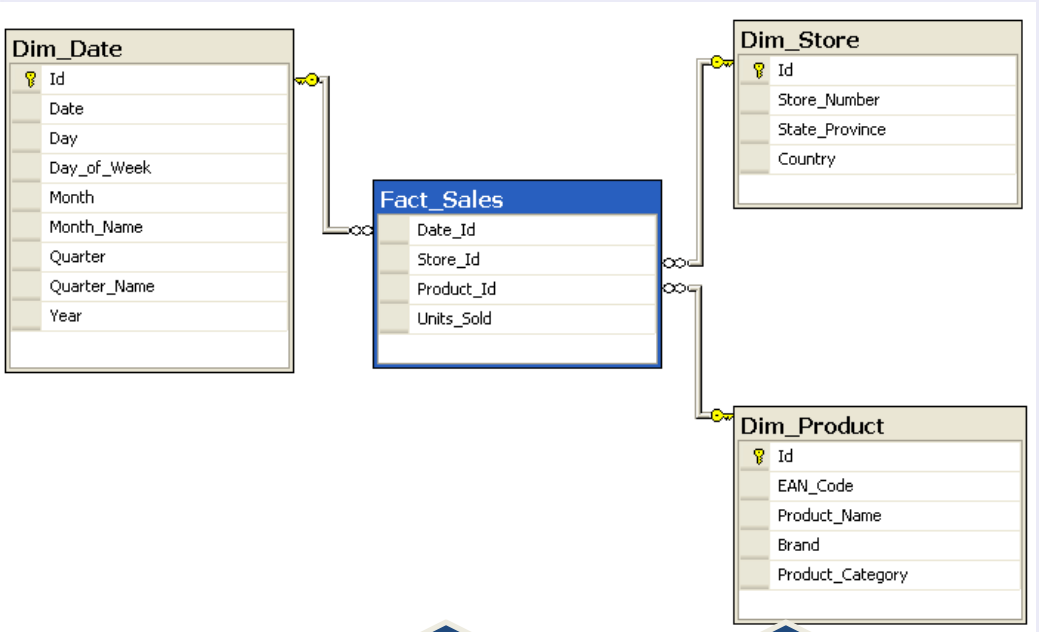

Users

Roles

Data

Model

Normal database architecture

OWASP 中国
The Open Web Application Security Project

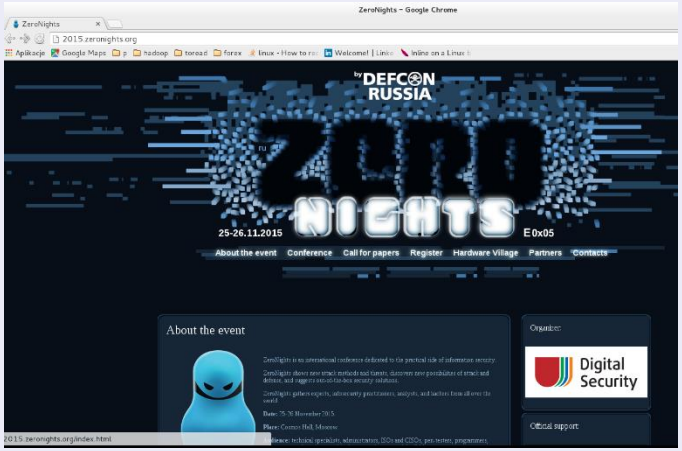

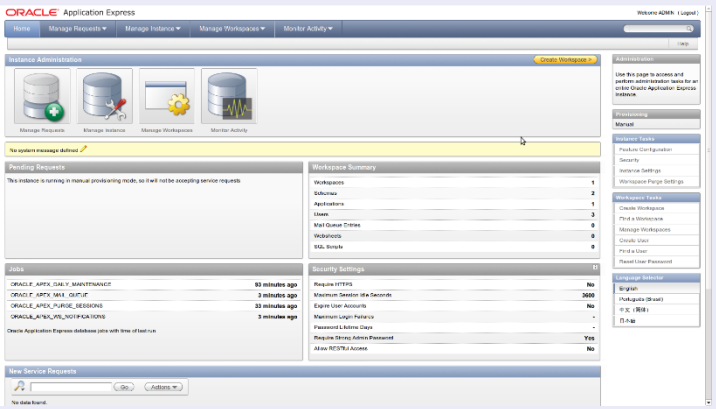

http://hackaday.com/2014/04/04/sql-injection-fools-speed-traps-and-clears-your-record/

Still normal database scenario

OWASP 中国
The Open Web Application Security Project

http://hackaday.com/2014/04/04/sql-injection-fools-speed-traps-and-clears-your-record/

http://nococonnect.blogspot.com/2015/06/red-light-cameras-in-columbia.html

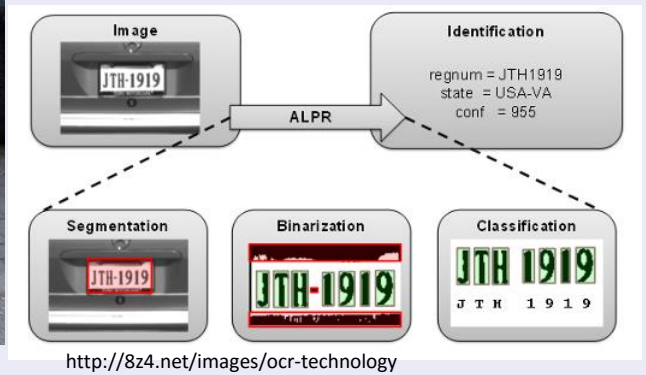

http://8z4.net/images/ocr-technology

CWE-xxx: SQL Injection through license plate

Normal database injection points

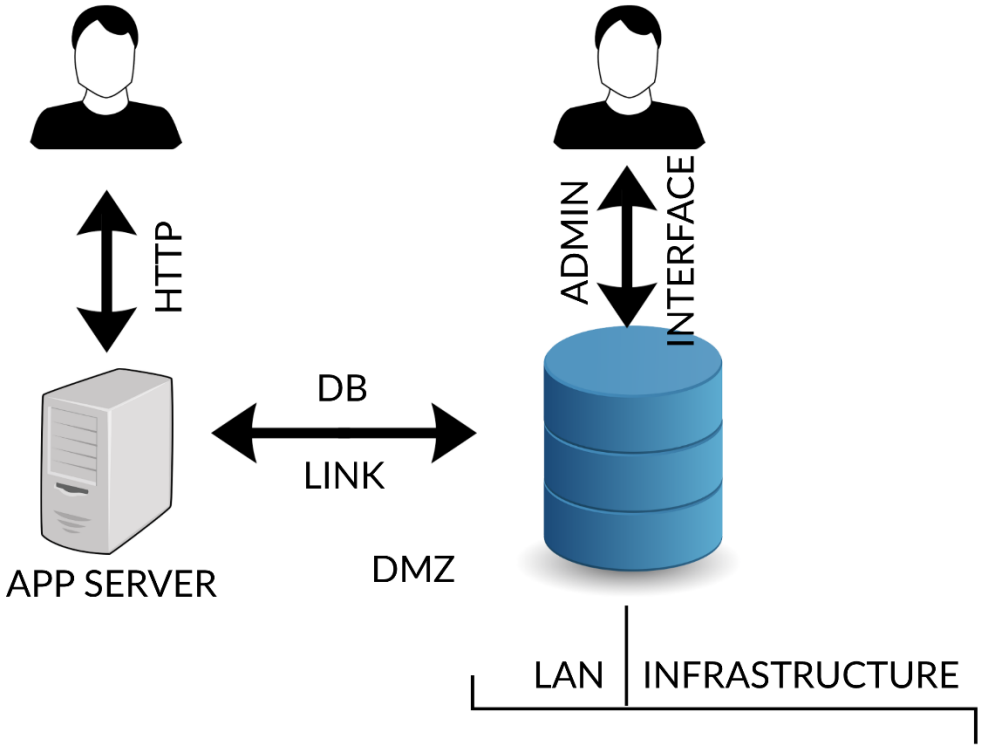

Normal database

Clear rules

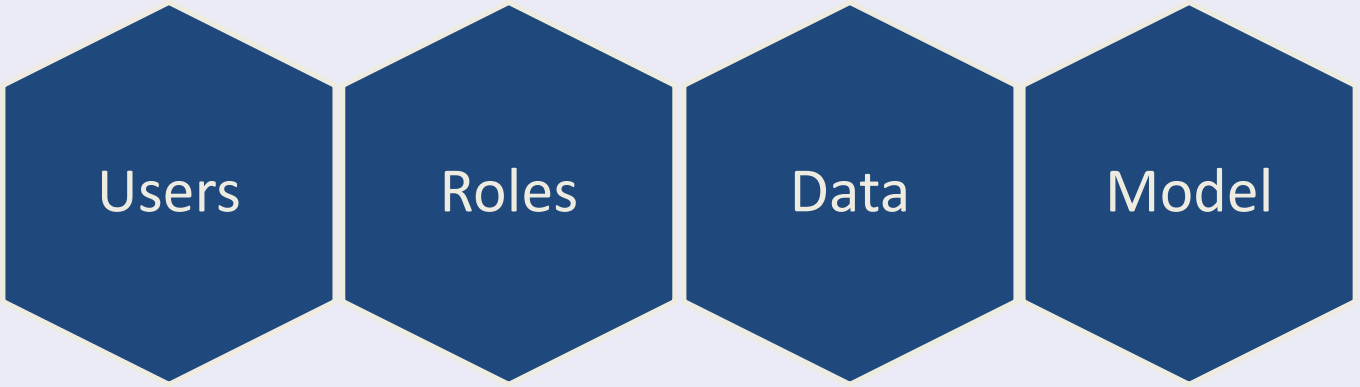

Clear target

Anecdote

OWASP 中国
The Open Web Application Security Project

user db, a lot of clients	critical banking data, one supplier

Only one common table

Q: Why don’t you split it into 2 dbs with a db link?
A: Too much effort and we want to have fast statistics from all data.

What is Hadoop?

http://fiveprime.org/blackmagic.cgi?id=7007203773

https://www.flickr.com/photos/photonquantique/2596581870/

Hadoop architecture schema

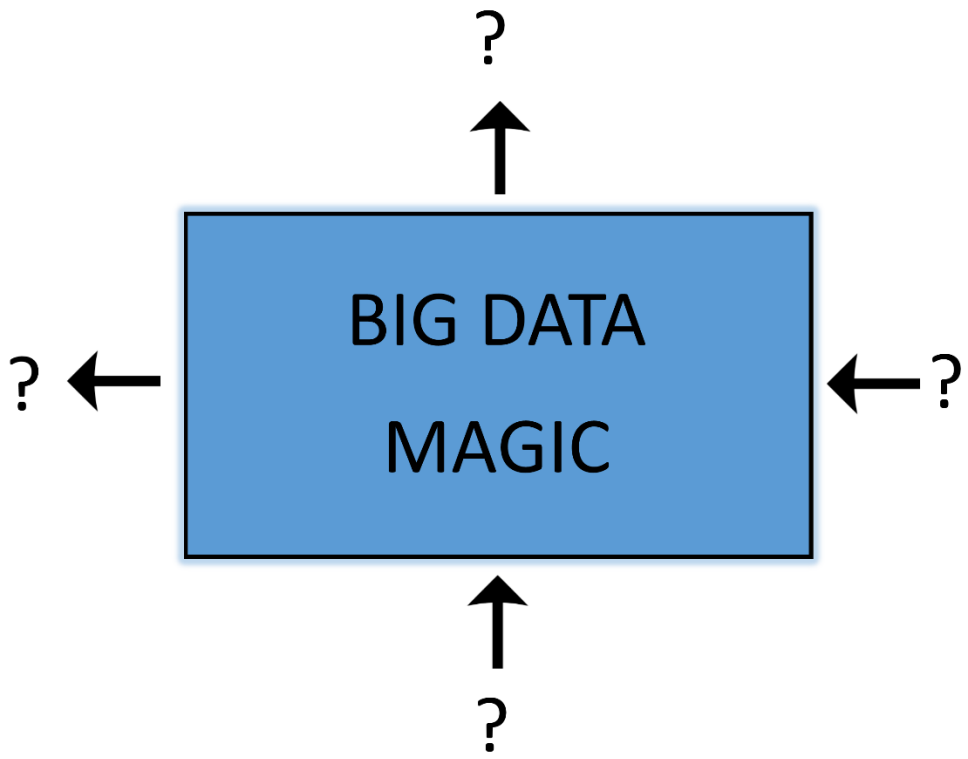

More on Hadoop

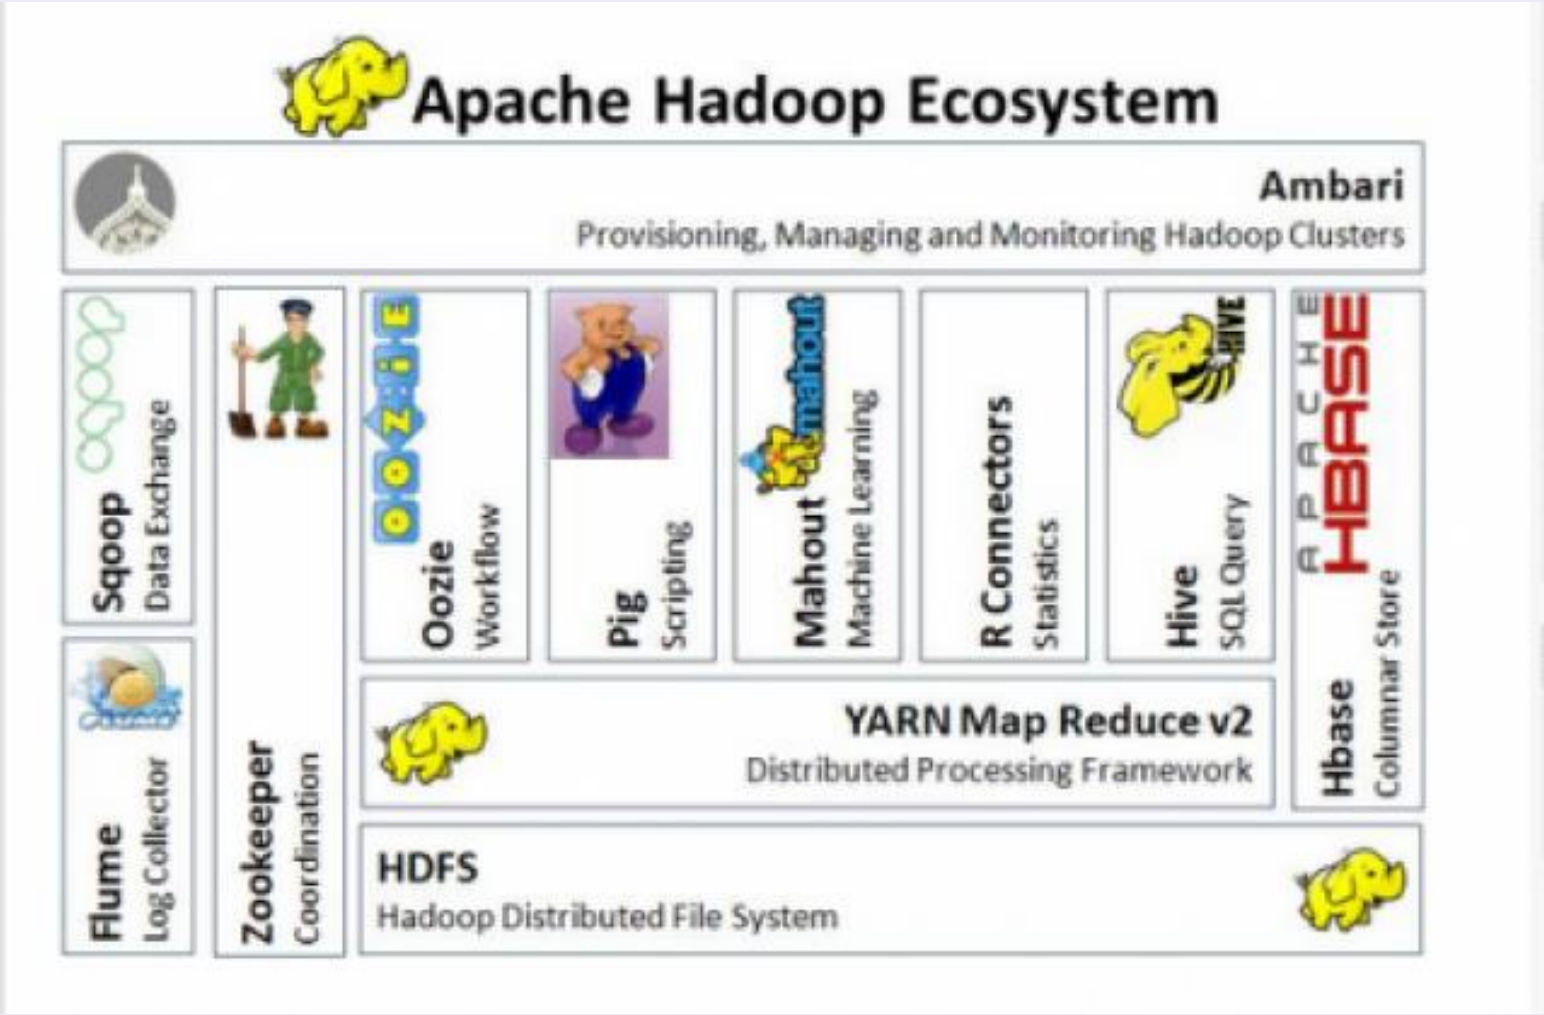

Hadoop injection points

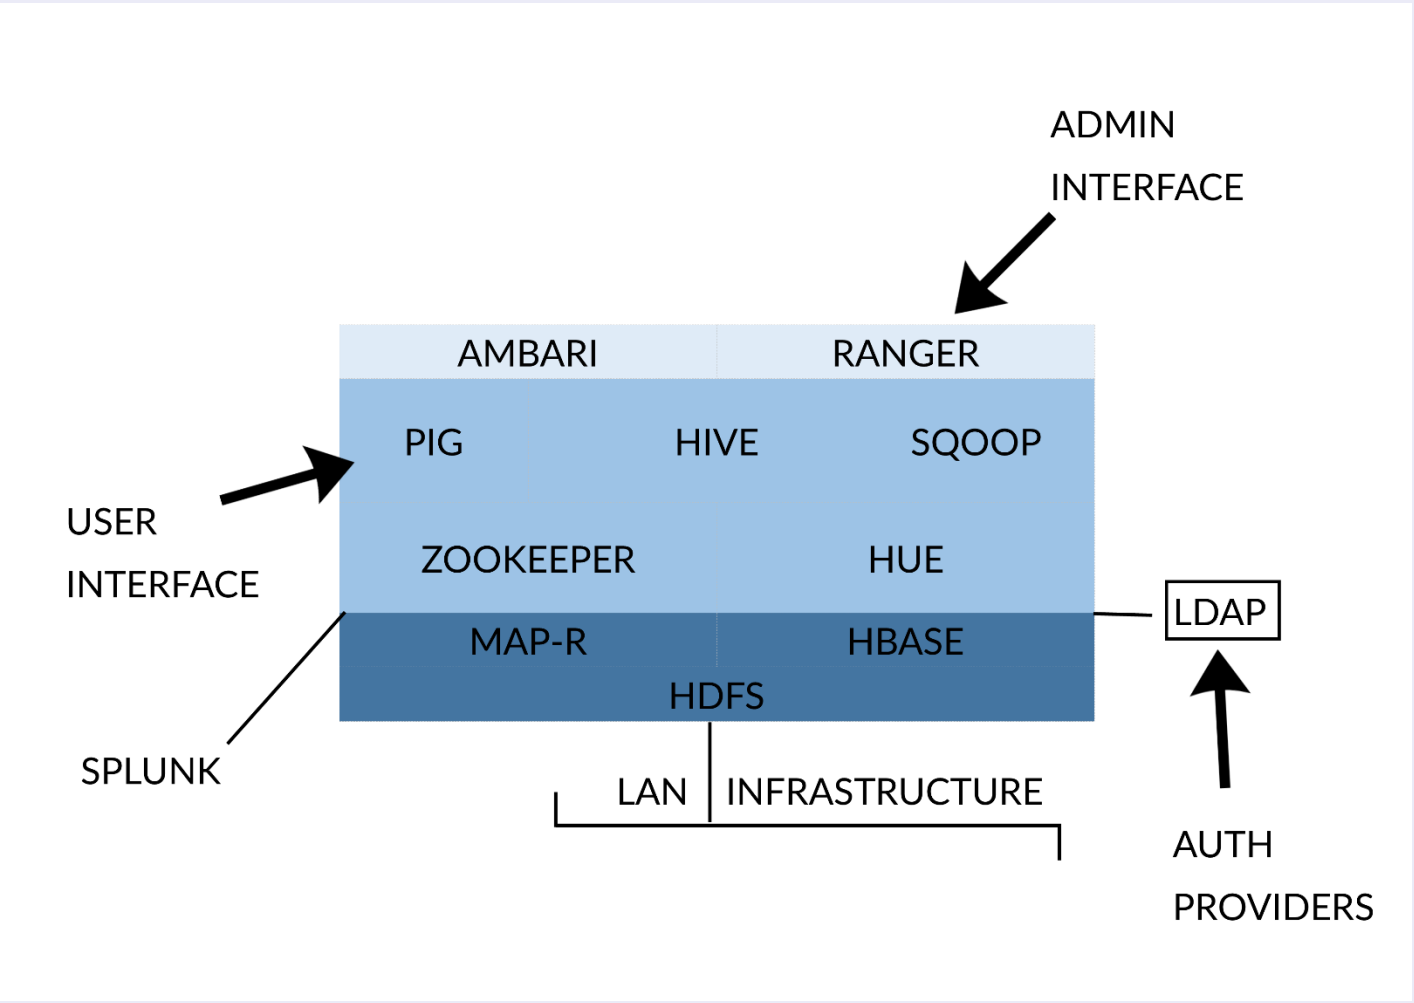

Hadoop scenario

OWASP 中国
The Open Web Application Security Project

https://www.flickr.com/photos/mattimattila/8349565473

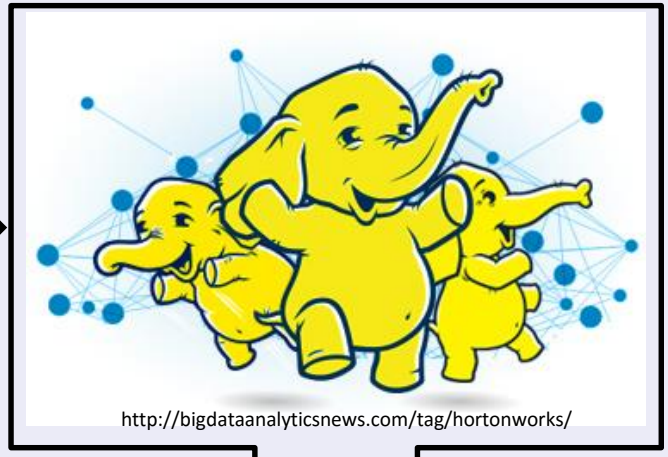

http://bigdataanalyticsnews.com/tag/hortonworks/

https://en.wikipedia.org/wiki/Moneygami

What is a lot of data?

OWASP 中国
The Open Web Application Security Project

facebook

- 21 PB of storage in a single HDFS cluster
- 2000 machines
- 12 TB per machine (a few machines have 24 TB each)
- 1200 machines with 8 cores each + 800 machines with 16 cores each
- 32 GB of RAM per machine
- 15 map-reduce tasks per machine

What is a lot of data?

- Our latest assessment:
- 32 machines, 8 cores each
- 24TB per machine
- 64 GB of RAM per machine
- Almost 1 PB disk space and 2TB of RAM

http://mrrobot.wikia.com/wiki/E_Corp

Attacker perspective

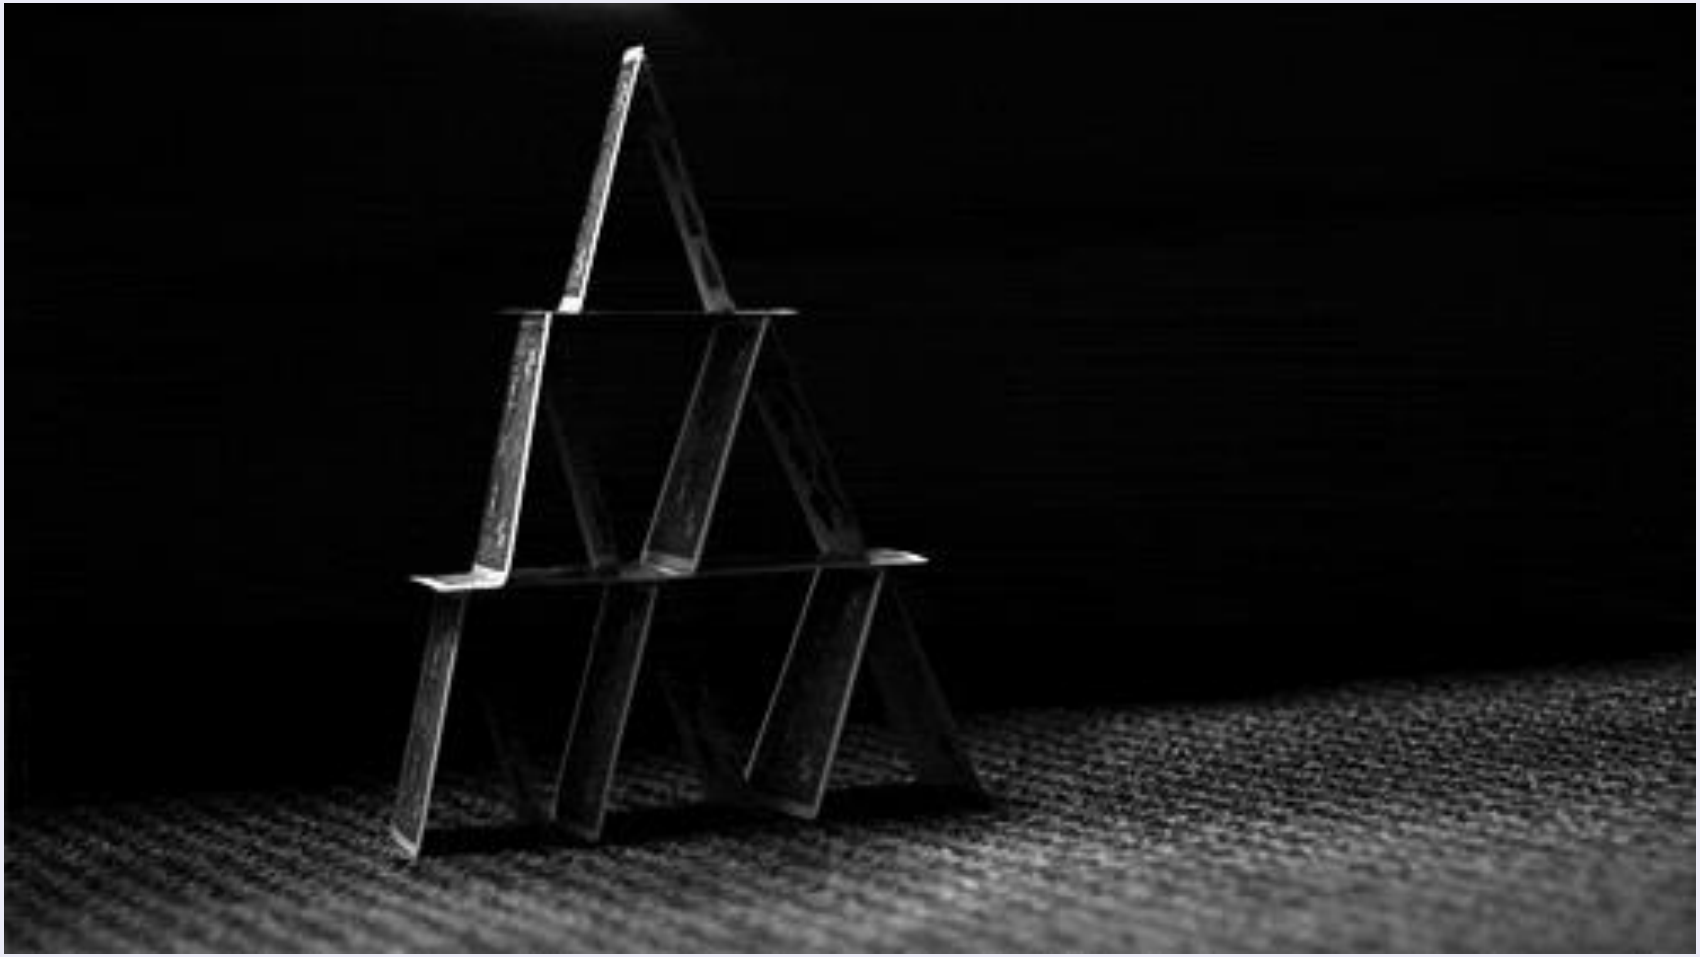

https://plus.google.com/+Magiccardtrickszonetips

OWASP 中国
The Open Web Application Security Project

Know your threats

RISK ANALYSIS

Risk analysis

OWASP 中国
The Open Web Application Security Project

Who → How → What

Who?

- Business perspective: competitor, script-kiddies, APT
- Technical perspective:

External attacker

- Anonymous
- Ex-employee

Insider

- Exployee (with some rights in Hadoop): user, admin
- Infected machine, APT

Risk analysis

Full compromise

Wojciech Dworakowski
IT Security Expert, Owner at SecuRing, OWASP Poland Chapter Leader

Following

Online banking owned by single attacker

Data safety vs. data security

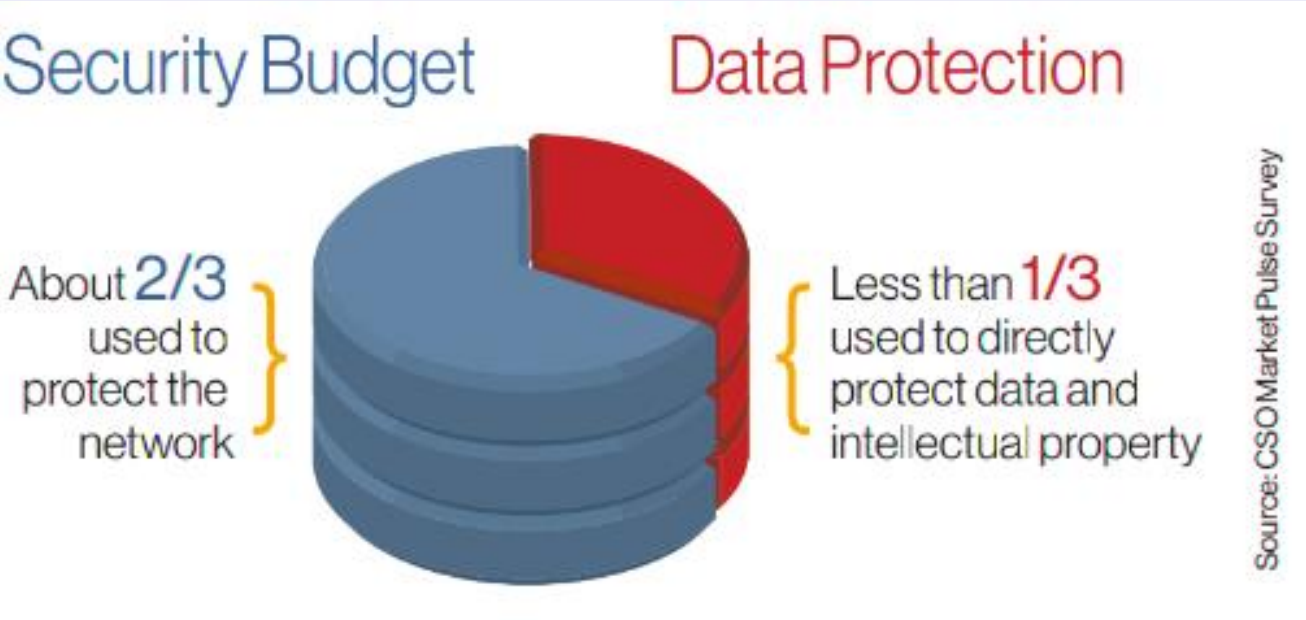

For what?

OWASP 中国
The Open Web Application Security Project

- *Q: What will be stored? A: „We do not know what data will be stored!”*
- Typical bank scenario
  - All transaction data
  - All sales data
  - All client data

https://www.reddit.com/r/gifs/comments/37aara/calculations_intensify/http://thewondrous.com/julia-gunthel-worlds

- Bigdata analytic says: „People who bought a dashcam are more likely to take a loan for a new car in the next month”

For what? Data theft

Forbes / Tech

FEB 16, 2012 @ 11:02 AM 2,866,944 VIEWS

How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did

MNN.com > Tech > Computers

How Facebook knows when you'll get divorced (even before you do)

Facebook knows who your romantic partner is, even if you keep that information private, and can even predict if the relationship will last.

Privilege escalation

- Authentication bypass

Abuse

- DoS
- Data tampering

Risk analysis

OWASP 中国
The Open Web Application Security Project

Who → How → What

How?

https://en.wikipedia.org/wiki/Dowsing#Rods

OWASP 中国
The Open Web Application Security Project

under sales-magic-cloud-big-data cover

WHAT HADOOP REALLY IS

Typical architecture

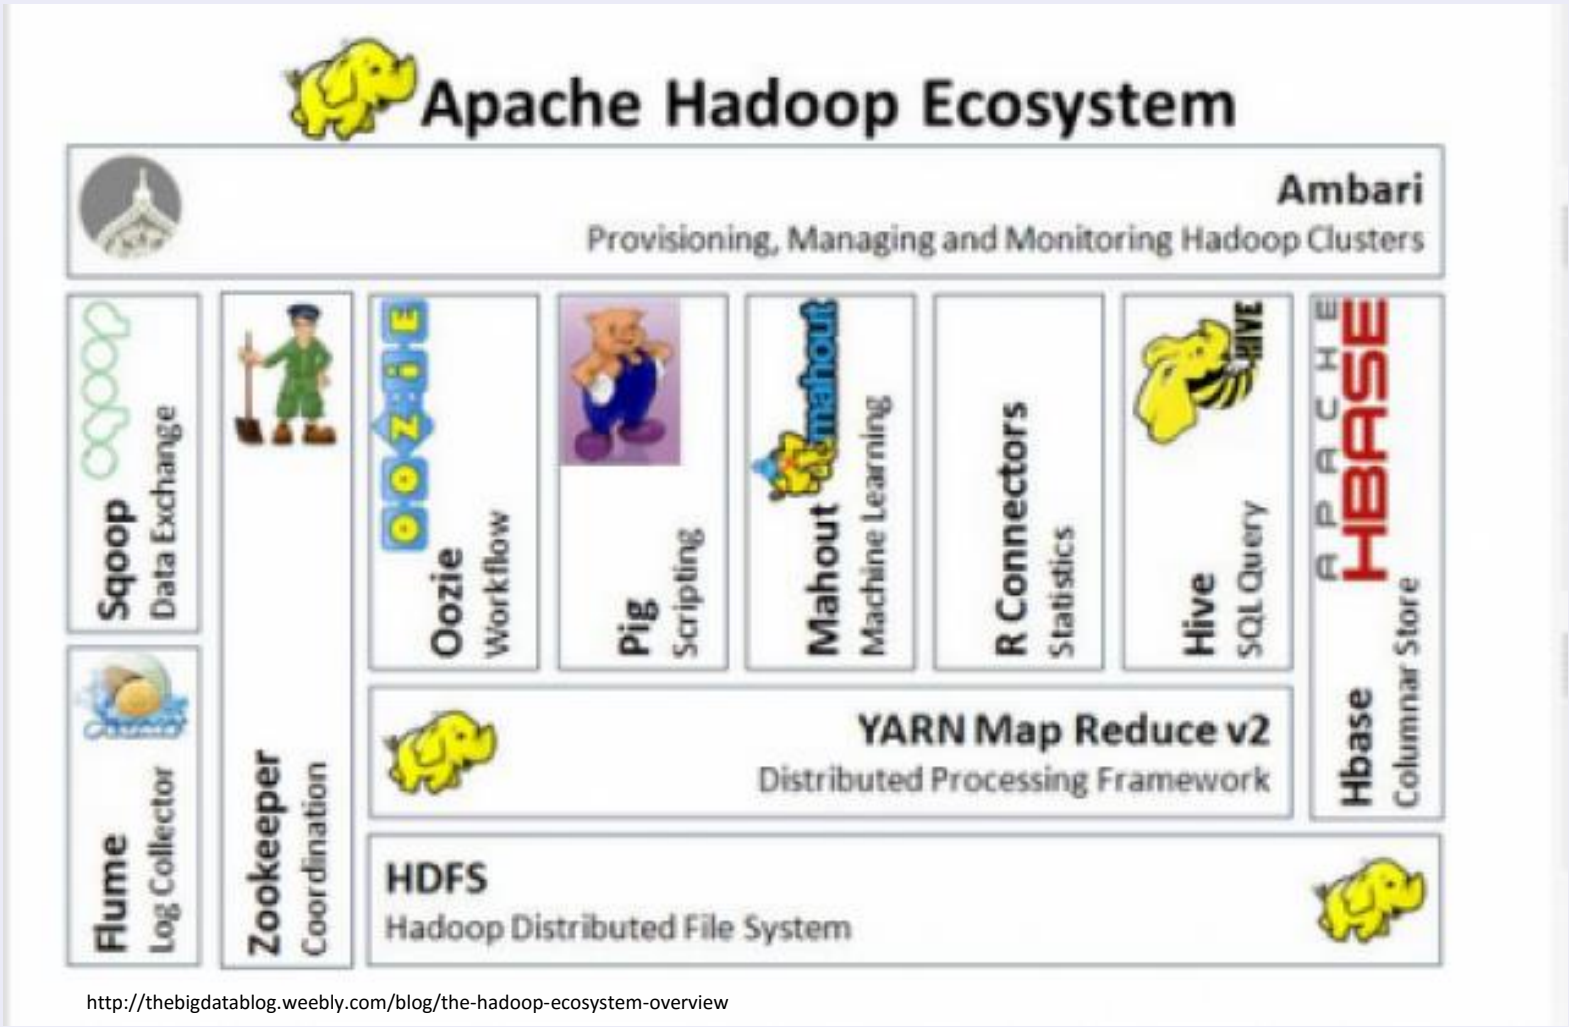

http://thebigdatablog.weebly.com/blog/the-hadoop-ecosystem-overview

Apache Hue

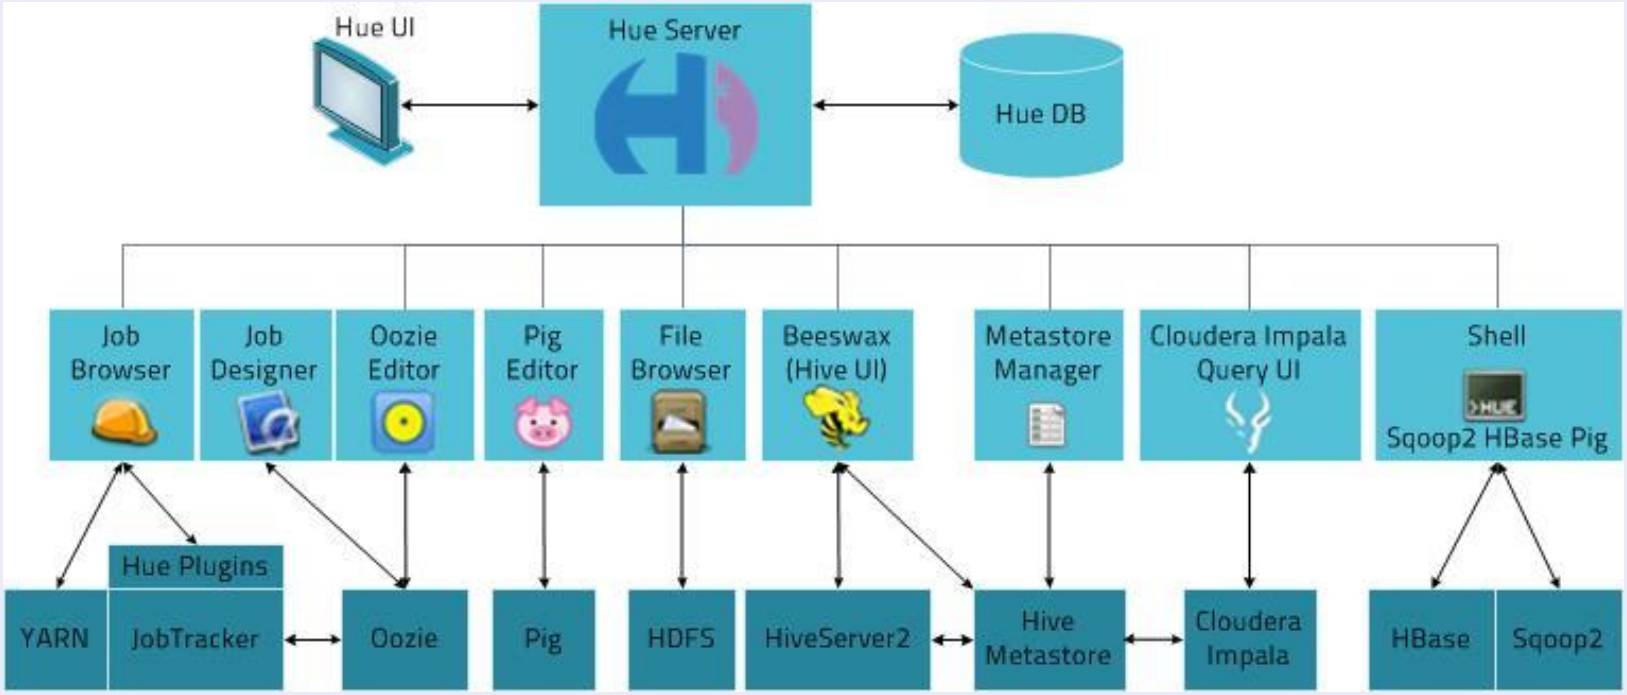

http://techbusinessintelligence.blogspot.com/2014/11/tableau-software-cloudera-hadoop.html

Hadoop injection points

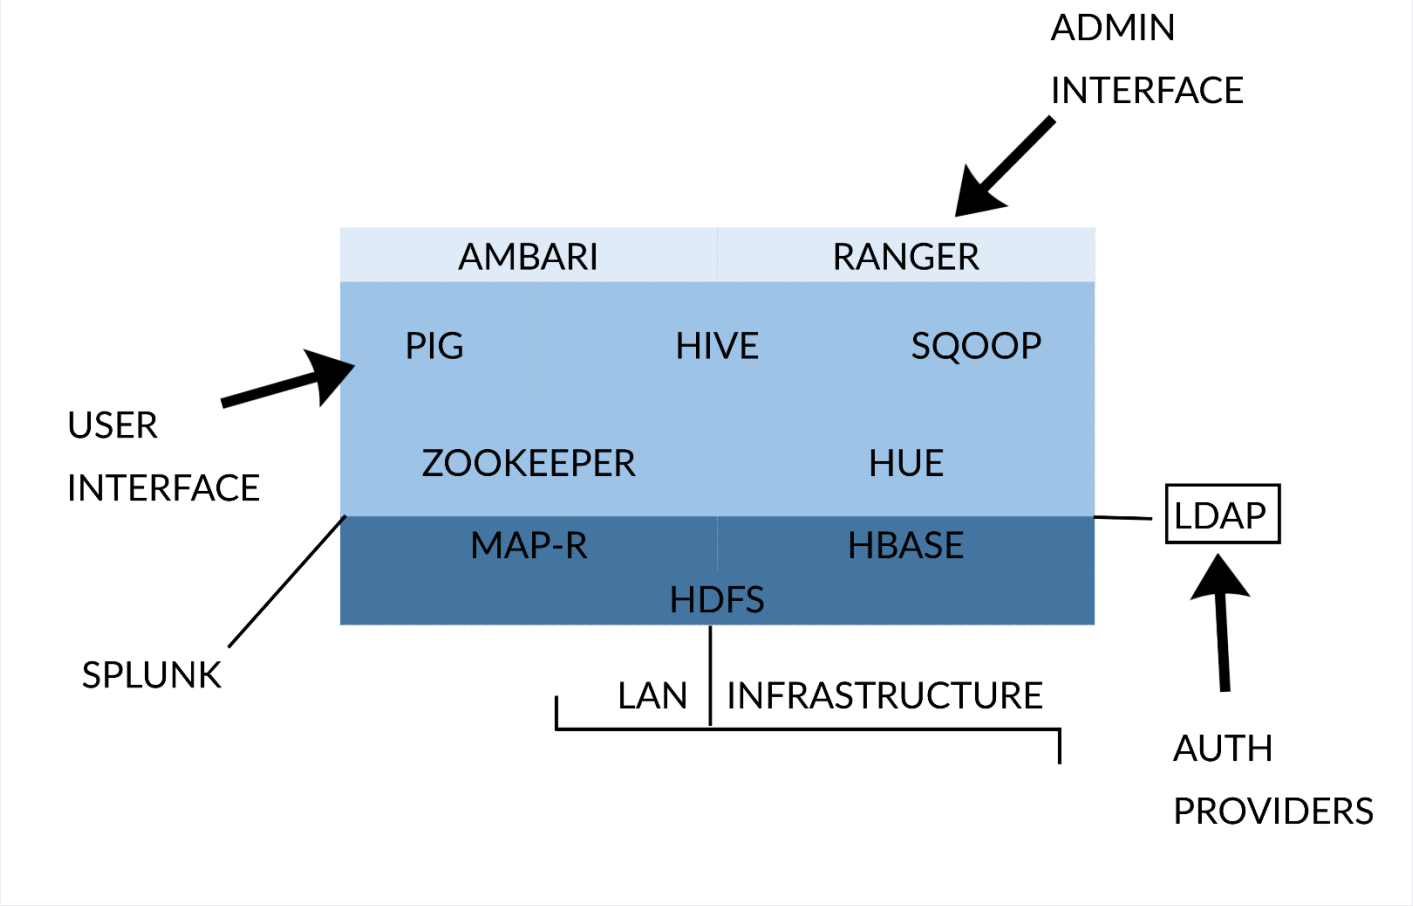

Differs much amongst distros

OWASP 中国
The Open Web Application Security Project

INTERFACES

Interfaces

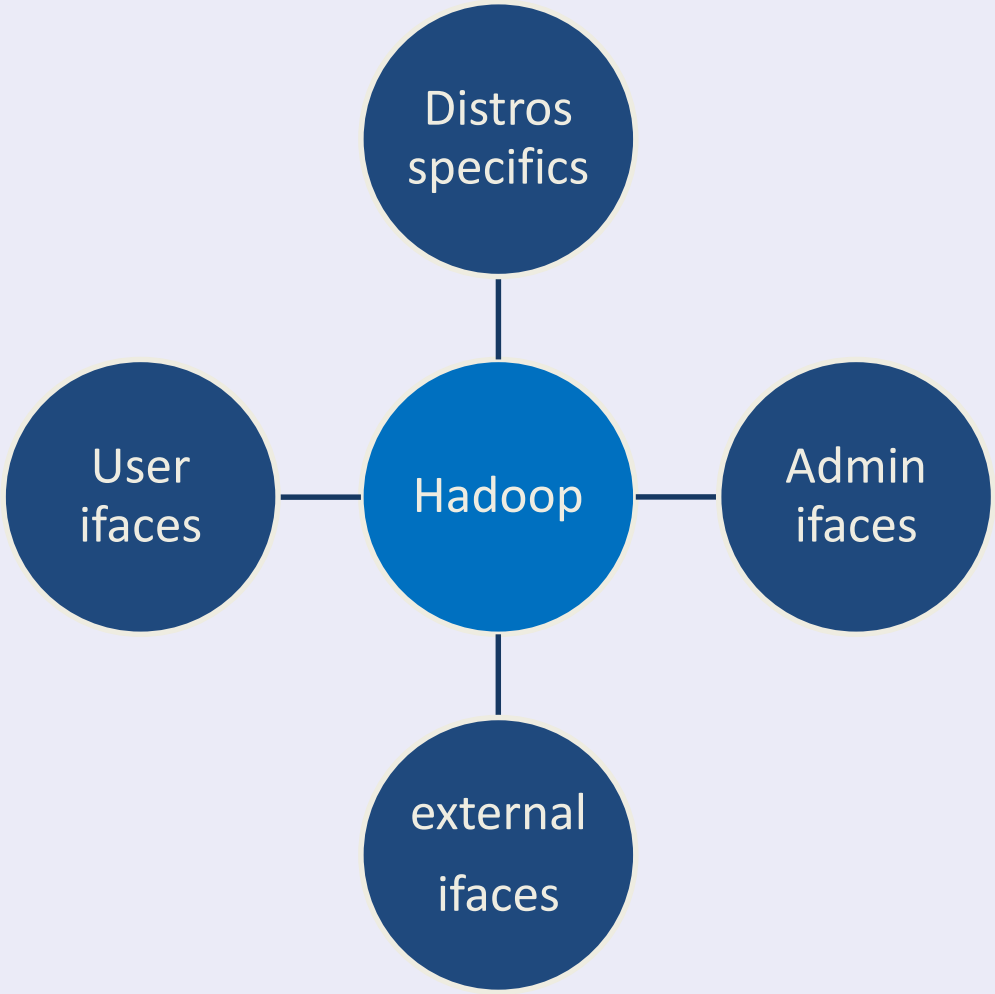
Distros specifics
User ifaces
Hadoop
Admin ifaces
external ifaces

OWASP 中国
The Open Web Application Security Project

a.k.a. crash course on hacking big data environments

OUR STORY WITH BIG DATA ASSESSMENT

Interfaces

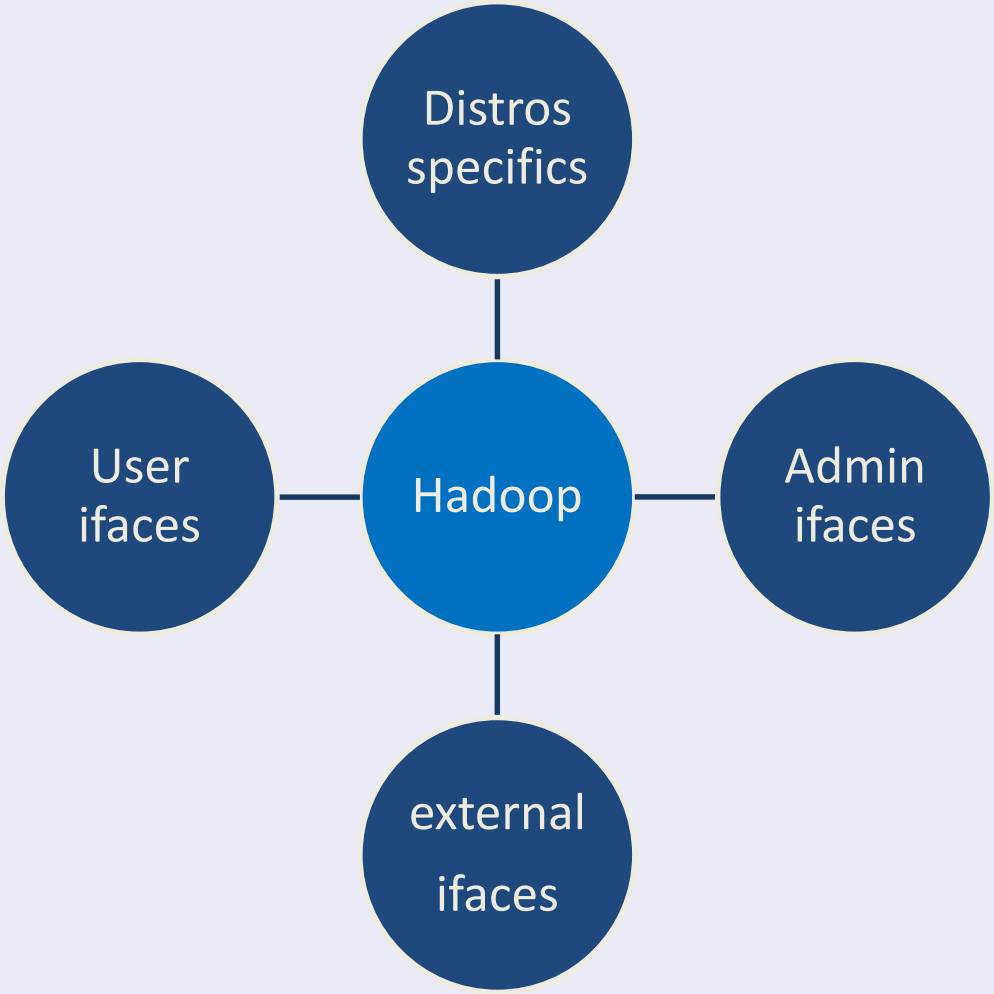
Distros specifics
User ifaces
Hadoop
Admin ifaces
external ifaces

OWASP 中国
The Open Web Application Security Project

for employees and applications

USER INTERFACES

User interfaces

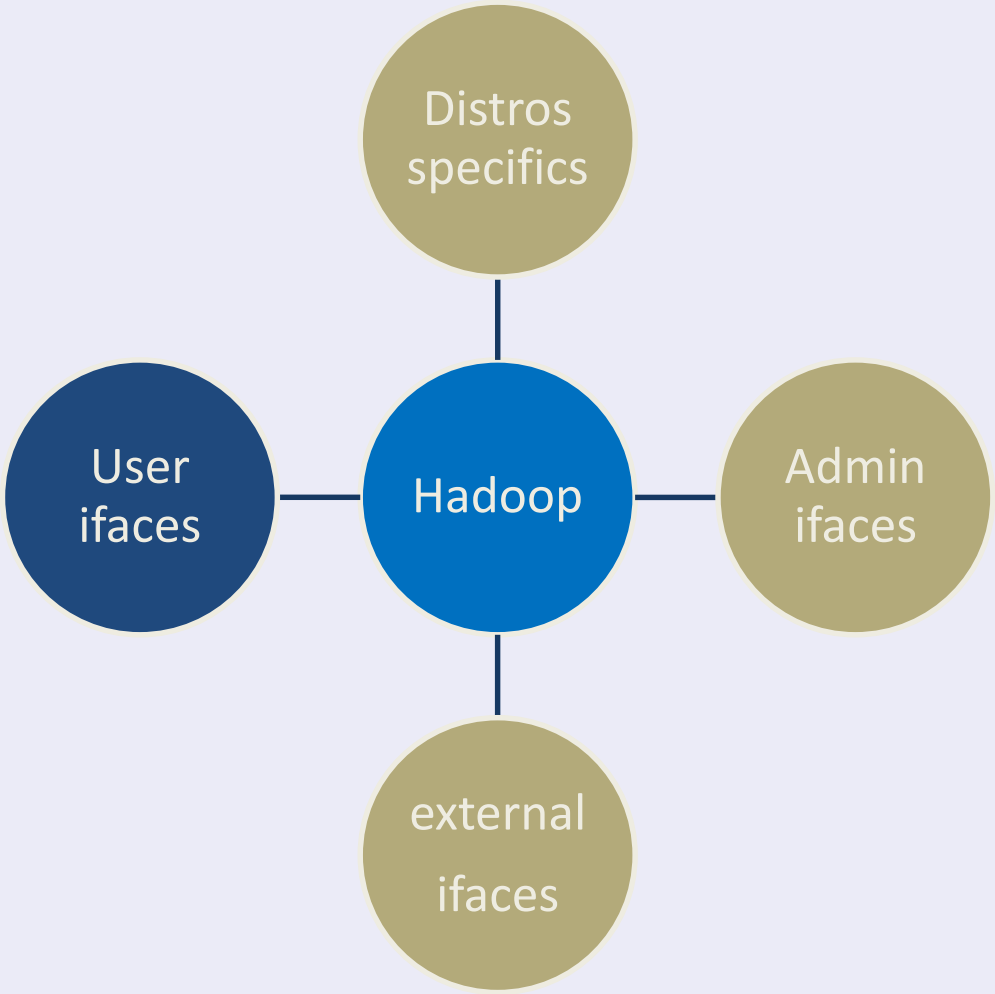

User interfaces

Apache Hue

- Pig, Hive, Impala, Hbase, Zookeeper, Mahout, Oozie

Other

- Tez, Solr, Slider, Spark, Phoenix, Accummulo, Storm

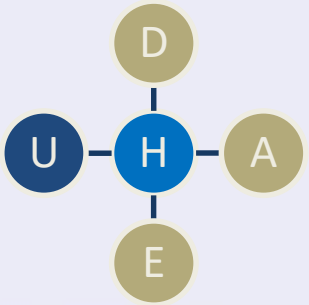

Is Hue an internal interface?

OWASP 中国
The Open Web Application Security Project

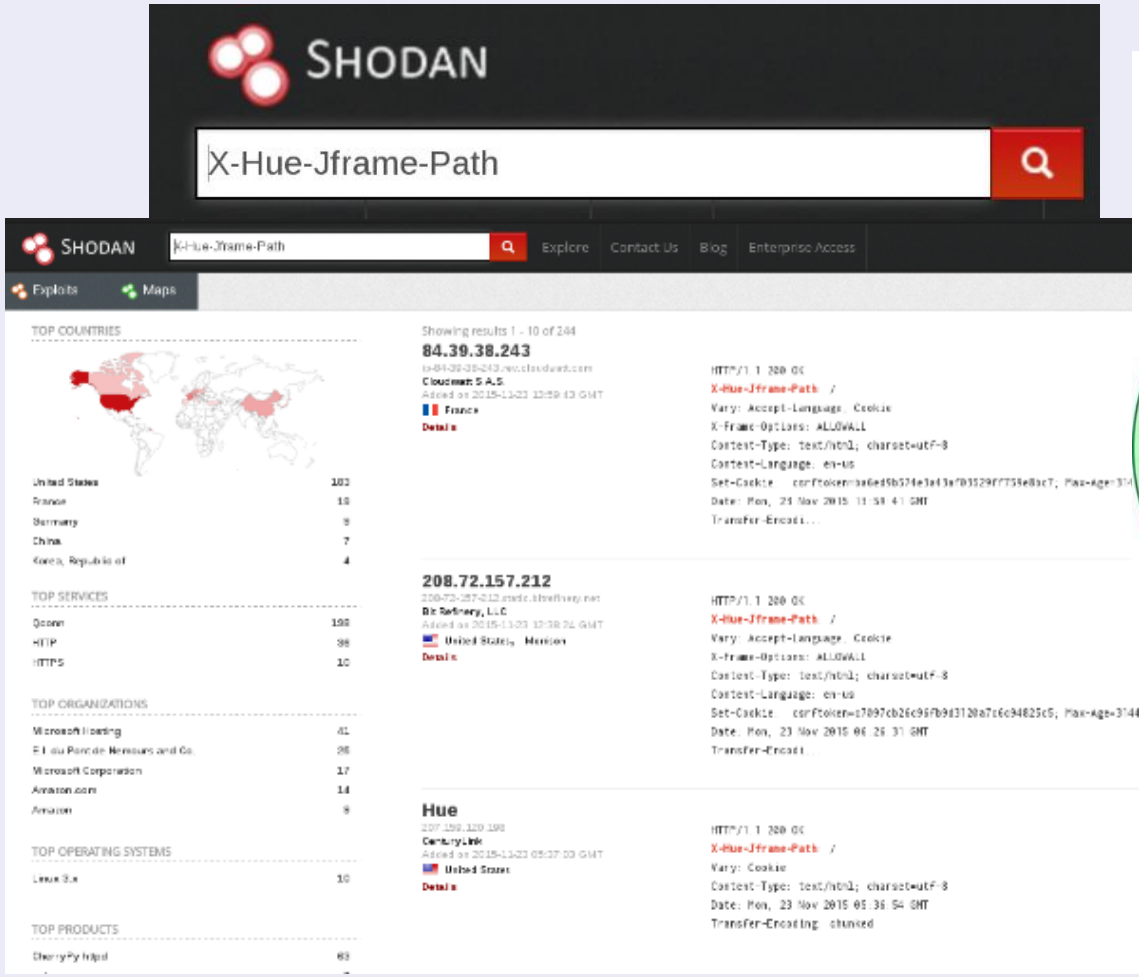

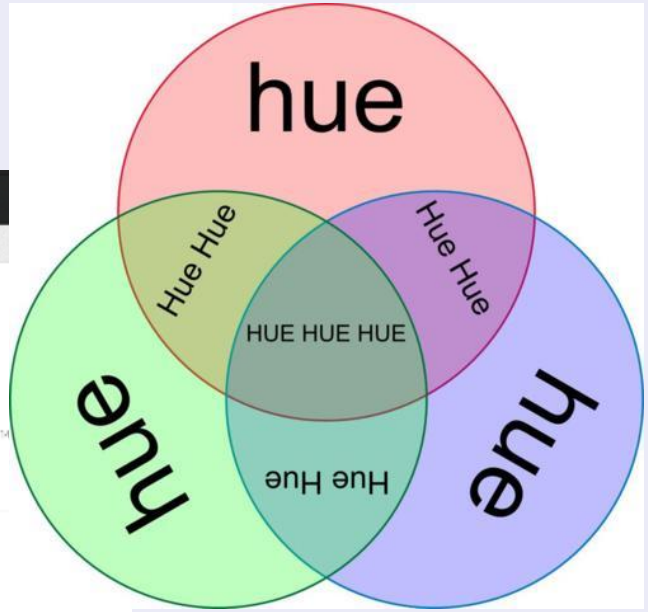

http://9gag.com/gag/awrwVL1/hue-hue-hue

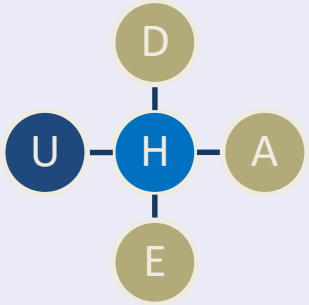

Apache Hue overview

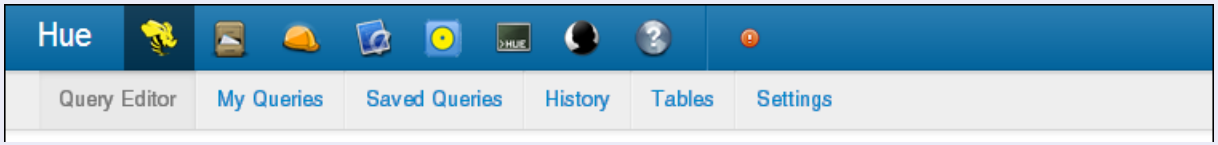

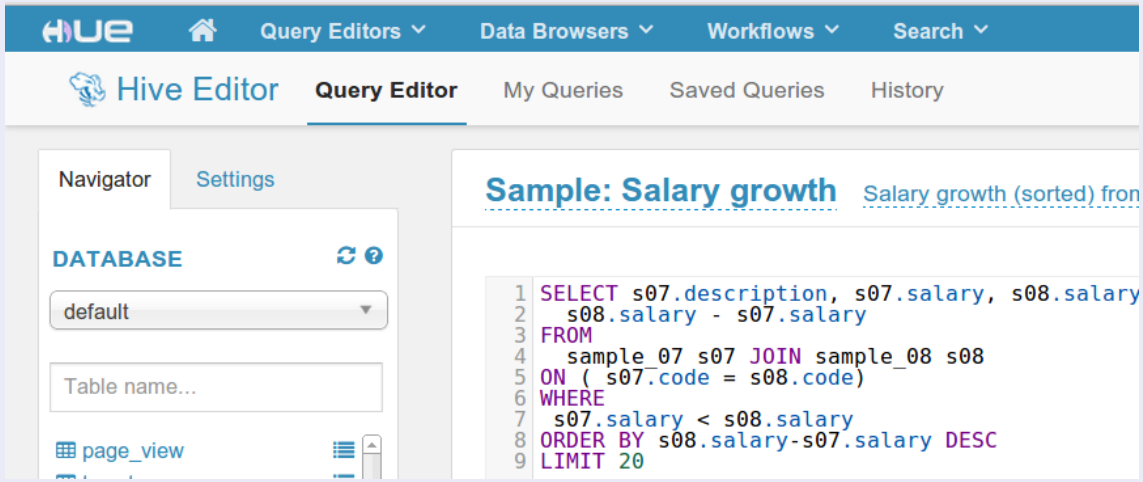

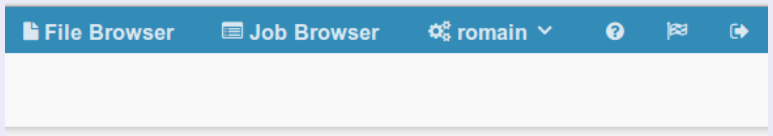

http://gethue.com/

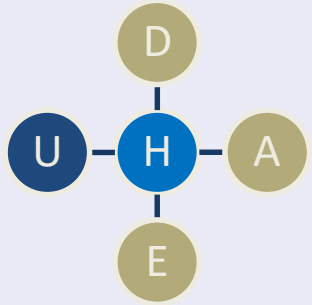

Apache Hue DOM XSS

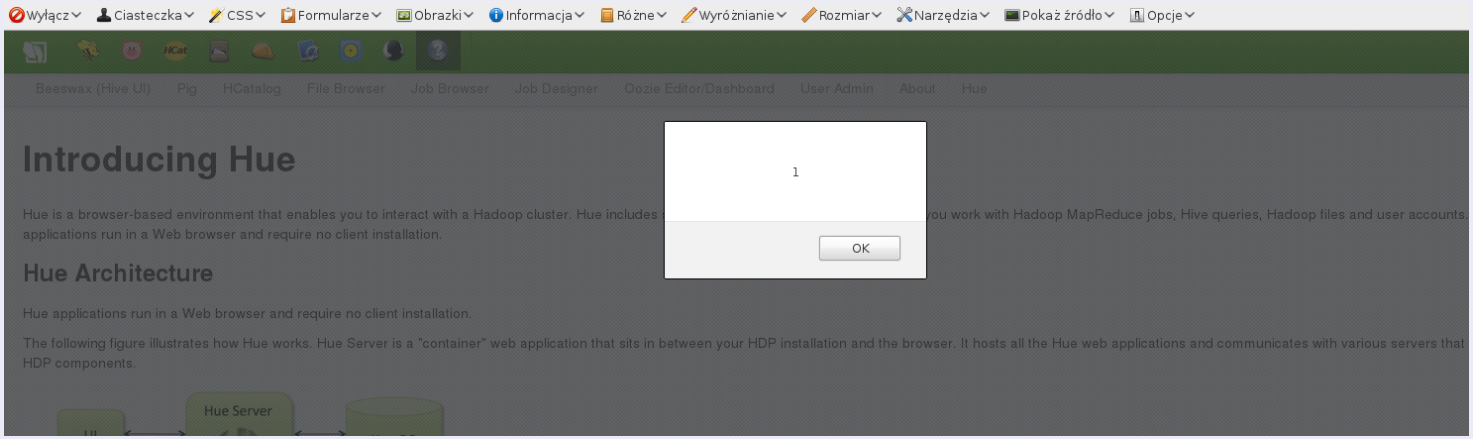

- `var _anchor = $("a[name='" + decodeURIComponent(window.location.hash.substring(1)) + "']").last();`
- **Payload: URL/help/#<img src="x" onerror="alert(1)">**

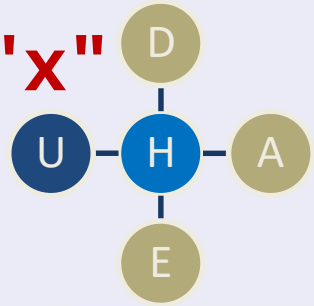

Apache Hue attack scenario

OWASP 中国
The Open Web Application Security Project

Target old Hadoop installation (with Hue 2.6.1, Django 1.2.3)

Target a user with access to Hue

Send him XSS

Get access to all Hadoop data designated for the user

Default configurations sucks

OWASP 中国
The Open Web Application Security Project

X-Frame-Options:ALLOWALL

SESSION_COOKIE_DOMAIN	None
SESSION_COOKIE_NAME	'sessionid'
SESSION_COOKIE_PATH	'/;HttpOnly'
SESSION_COOKIE_SECURE	True
SESSION_ENGINE	'django.contrib.sessions.backends.db'
SESSION_EXPIRE_AT_BROWSER_CLOSE	False
SESSION_FILE_PATH	None
SESSION_SAVE_EVERY_REQUEST	False
SETTINGS_MODULE	'desktop.settings'
SHORT_DATETIME_FORMAT	'm/d/Y P'
SHORT_DATE_FORMAT	'm/d/Y'
SITE_ID	1
SKIP_SOUTH_TESTS	True
TEMPLATE_CONTEXT_PROCESSORS	('django.contrib.auth.context_processors.auth', 'django.core.context_processors.debug', 'django.core.context_processors.media', 'django.contrib.messages.context_processors.mess
TEMPLATE_DEBUG	False
TEMPLATE_DIRS	('/usr/lib/hue/desktop/core/templates',)
TEMPLATE_LOADERS	('django.template.loaders.filesystem.load_template_source', 'desktop.lib.template_loader
TEMPLATE_STRING_IF_INVALID	''
TEST_DATABASE_CHARSET	None
TEST_DATABASE_COLLATION	None
TEST_DATABASE_NAME	None
TEST_RUNNER	'django.test.simple.DjangoTestSuiteRunner'
THOUSAND_SEPARATOR	','
TIME_FORMAT	'P'
TIME_INPUT_FORMATS	('%H:%M:%S', '%H:%M')
TIME_ZONE	'America/Los_Angeles'
TRANSACTIONS_MANAGED	False
URL_VALIDATOR_USER_AGENT	'Django/1.2.3 (http://www.djangoproject.com)'
USE_ETAGS	False
USE_I18N	True
USE_L10N	True
USE_THOUSAND_SEPARATOR	False
X_FRAME_OPTIONS	'ALLOWALL'
YEAR_MONTH_FORMAT	'F Y'

You're seeing this error because you ha e DEBUG = True in your Django settings file. Change that to False, and Django will display a standard 500 page.

debug 7 of 7

Page not found (404)

Request Method: GET
Request URL: https://lxhulo001:8000/x

Using the URLconf defined in desktop.urls, Django tried these URL p

1. ^about/static/(?P<path>.*)$
2. ^beeswax/static/(?P<path>.*)$
3. ^filebrowser/static/(?P<path>.*)$
4. ^hcatalog/static/(?P<path>.*)$
5. ^help/static/(?P<path>.*)$
6. ^jobbrowser/static/(?P<path>.*)$
7. ^jobsub/static/(?P<path>.*)$
8. ^oozie/static/(?P<path>.*)$
9. ^pig/static/(?P<path>.*)$
10. ^shell/static/(?P<path>.*)$
11. ^useradmin/static/(?P<path>.*)$
12. ^static/(?P<path>.*)$
13. ^(?P<path>favicon.ico)$
14. ^accounts/login/$
15. ^accounts/logout/$
16. ^logs$
17. ^dump_config$
18. ^download_logs$
19. ^bootstrap.js$
20. ^profile$
21. ^prefs/(?P<key>\w+)?$
22. ^status_bar/?$
23. ^admin/
24. ^debug/threads$
25. ^debug/who_am_i$
26. ^debug/check_config$
27. ^debug/check_config_ajax$
28. ^log_frontend_event$
29. ^jasmine
30. ^$
31. ^about/
32. ^beeswax/
33. ^filebrowser/
34. ^hcatalog/
35. ^help/
36. ^jobbrowser/
37. ^jobsub/
38. ^oozie/
39. ^pig/

OWASP 中国
The Open Web Application Security Project

for admins and maintenance

ADMIN INTERFACES

Admin interfaces

OWASP 中国
The Open Web Application Security Project

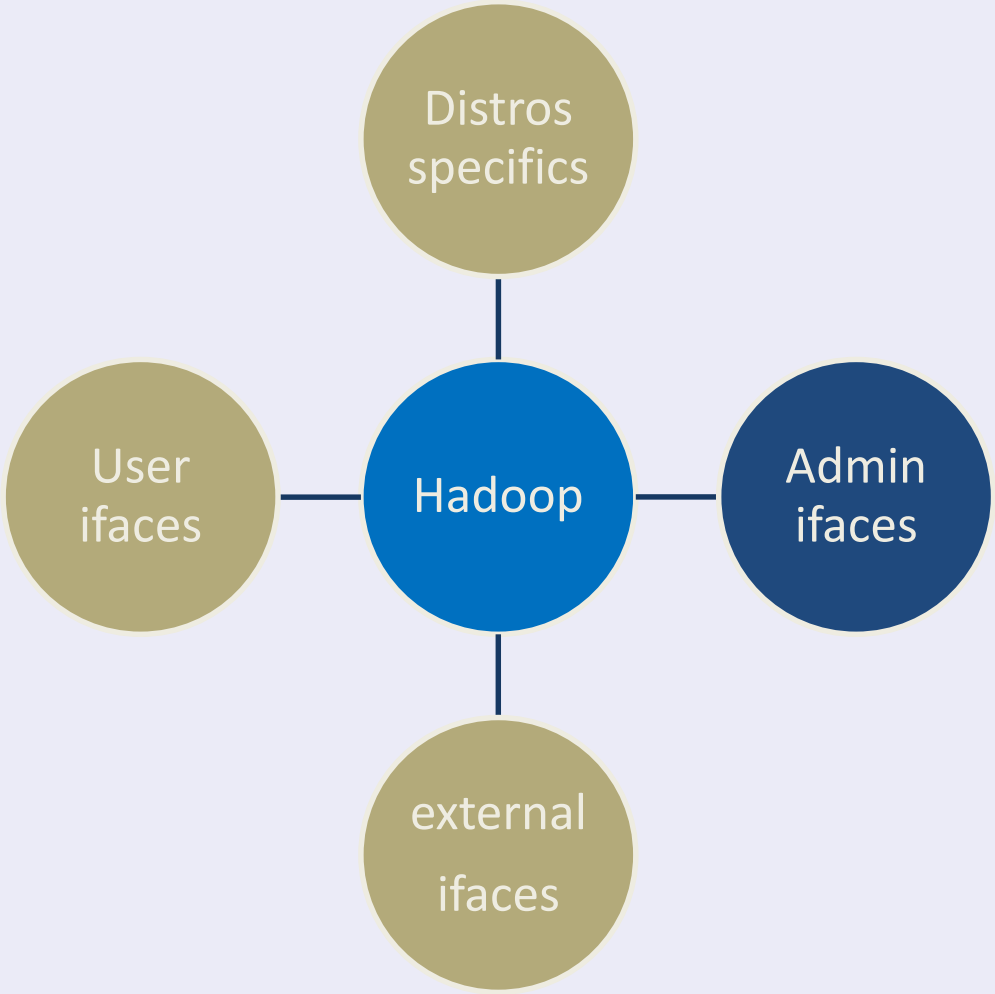

Admin interfaces

Apache Ambari

- Provisioning, monitoring

Apache Ranger

- Security: authorization, authentication, auditing, data encryption, administration

Other

- Knox, Cloudbreak, Zookeeper, Falcon, Atlas, Sqoop, Flume, Kafka

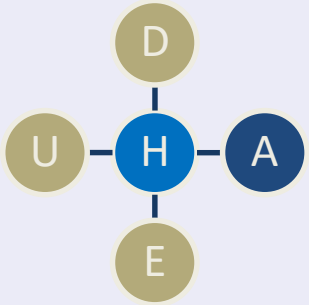

Apache Ambari

- About Ambari

Feature	Benefit
Wizard-driven interface	Facilitates installation of Hadoop across any number of hosts
API-driven installations	Ambari Blueprints for automated provisioning
Granular service control	Precise management of Hadoop services and component lifecycles
Configuration change history	Ongoing management of Hadoop service configurations
RESTful APIs	Enables integration with enterprise systems
Extensible framework	Brings custom services under management via Ambari Stacks
Customizable user interface	Develop innovative user experiences via Ambari Views Framework
User Views	Advanced capabilities for cluster optimization and tuning for Hadoop DevOps

http://www.slideshare.net/hortonworks/ambari-using-a-local-repository?next_slideshow=1

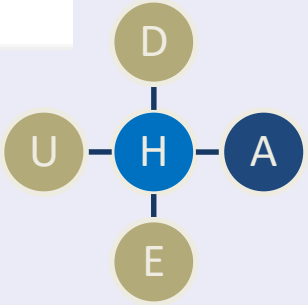

Apache Ambari

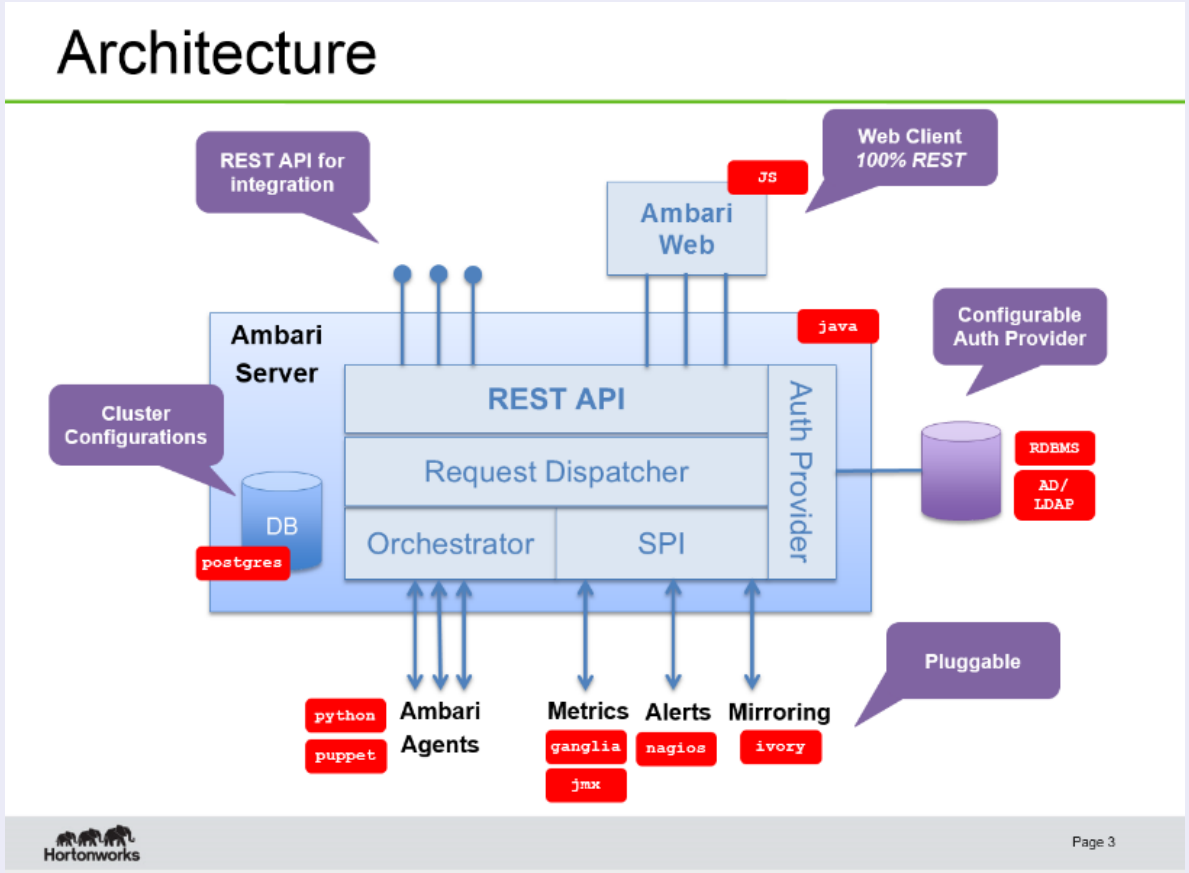

http://www.slideshare.net/hortonworks/ambari-using-a-local-repository?next_slideshow=1

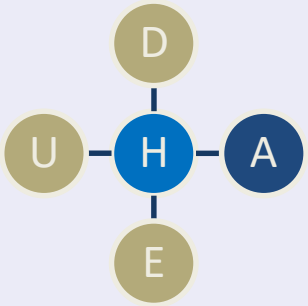

Is Ambari an internal interface?

OWASP 中国
The Open Web Application Security Project

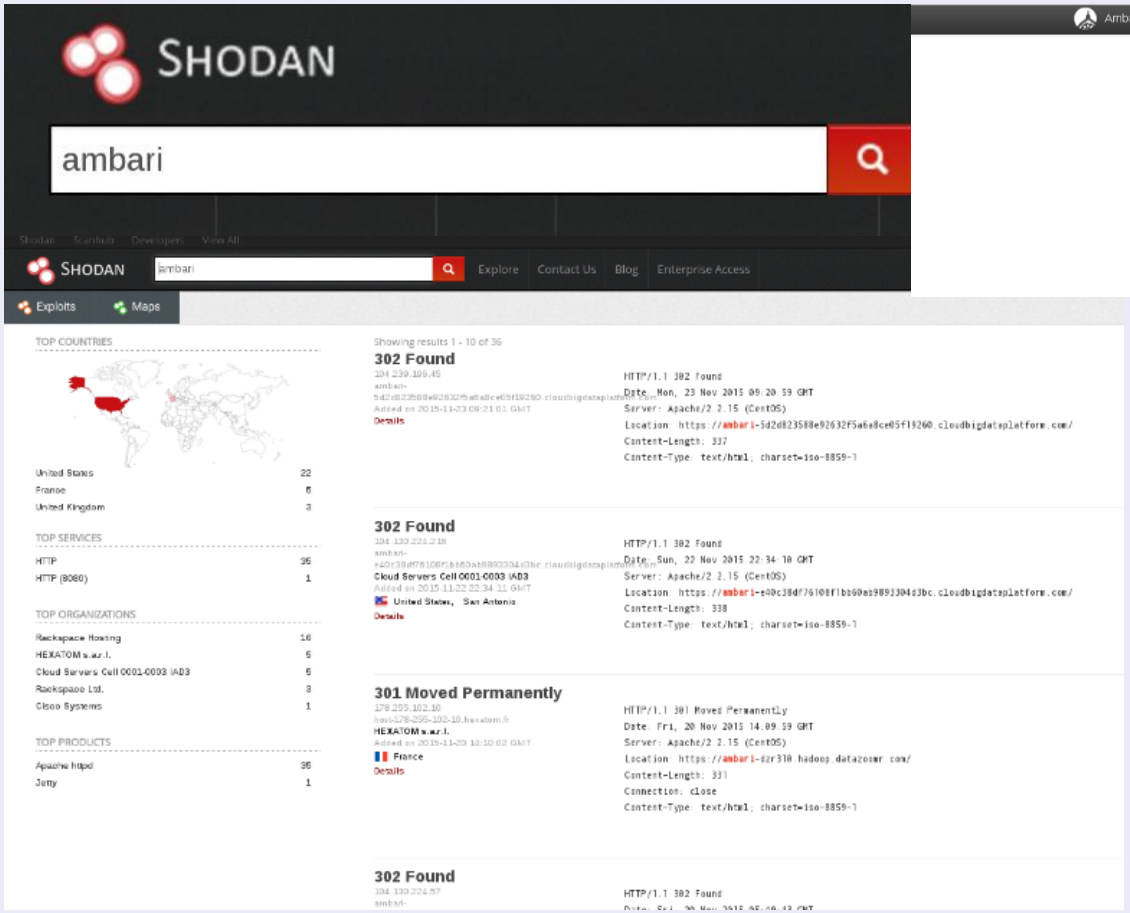

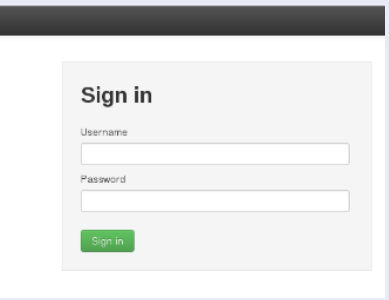

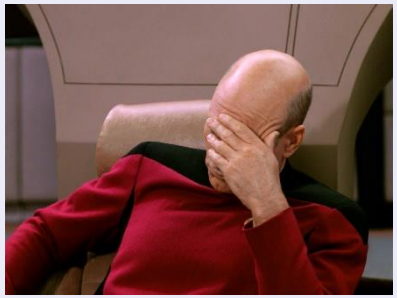

http://knowyourmeme.com/memes/facepalm

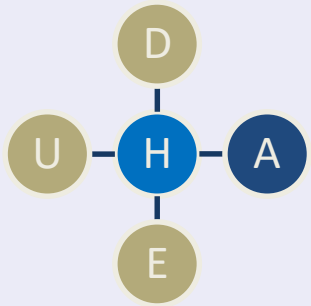

Apache Ambari

- Standard users can sign into Ambari (WHY?)
- Low hanging fruits: directory listing by default, no cookie flags, no CSRF protection
- Interesting proxy script ->

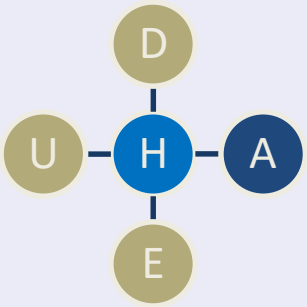

Apache Ambari REST API proxy

OWASP 中国
The Open Web Application Security Project

Standard request:

```
/proxy?url=http://XXXXXXXXX:8188/ws/v1/ti
meline/HIVE_QUERY_ID?limit=1&secondaryFil
ter=tez:true&_=1424180016625
```

Tampered request (logs accessible only from DMZ):

```
/proxy?url=http://google.com
/proxy?url=http://XXXXXXX:8088/logs
/proxy?url=http://XXXXXXX:8088/logs
/yarn-yarn-resourcemanager-
XXXXXXX.log
```

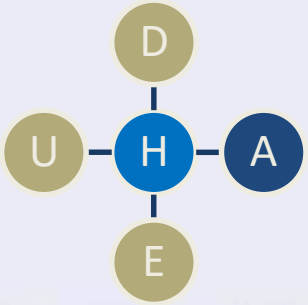

Apache Ambari Server Side Request Forgery

OWASP 中国
The Open Web Application Security Project

Directory: /logs/

hadoop-mapreduce.jobsummary.log		137797 bytes	Jan 22, 2015 6:18:54 PM
yarn-yarn-historyserver-	.log	3866624 bytes	Feb 16, 2015 11:23:02 AM
yarn-yarn-historyserver-	.out	4096 bytes	Feb 14, 2015 2:08:00 PM
yarn-yarn-historyserver-	.out.1	828 bytes	Dec 10, 2014 11:51:13 AM
yarn-yarn-historyserver-	.out.2	828 bytes	Dec 10, 2014 11:44:31 AM
yarn-yarn-historyserver-	.out.3	828 bytes	Dec 10, 2014 10:55:43 AM
yarn-yarn-resourcemanager-	.log	19779584 bytes	Feb 16, 2015 11:24:22 AM
yarn-yarn-resourcemanager-	.out	171856 bytes	Feb 15, 2015 1:25:50 PM
yarn-yarn-resourcemanager-	.out.1	2192 bytes	Dec 10, 2014 12:46:05 PM
yarn-yarn-resourcemanager-	.out.2	2086 bytes	Dec 10, 2014 11:46:30 AM
yarn-yarn-resourcemanager-	.out.3	2086 bytes	Dec 10, 2014 11:00:48 AM

CVE-2015-1775

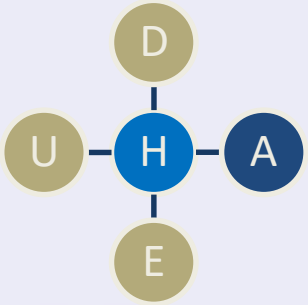

Apache Ambari attack scenario

- Target old Hadoop installation with Ambari 1.5.0 to 2.0.2
- Hijack standard account (or use Hue XSS to perform CSRF)
- Log into Ambari, use CVE-2015-1775
- Get access to local network (DMZ) – HTTP only
- Download logs, exploit other Hadoop servers in DMZ

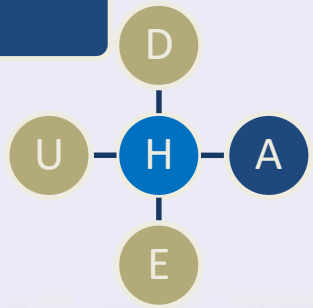

Admin interfaces

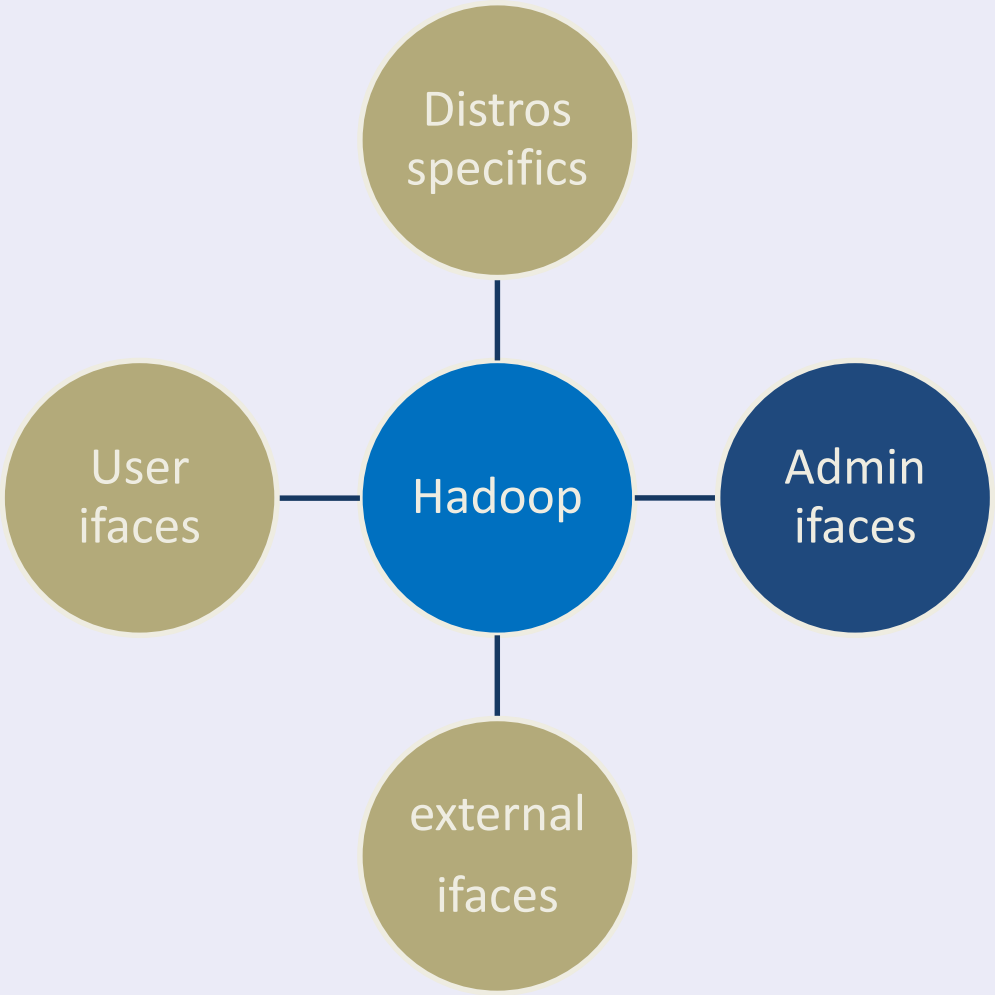

Apache Ranger overview

- Previously: Apache Argus, XA-Secure
- Provides central administration for policies, users/groups, analytics and audit data.

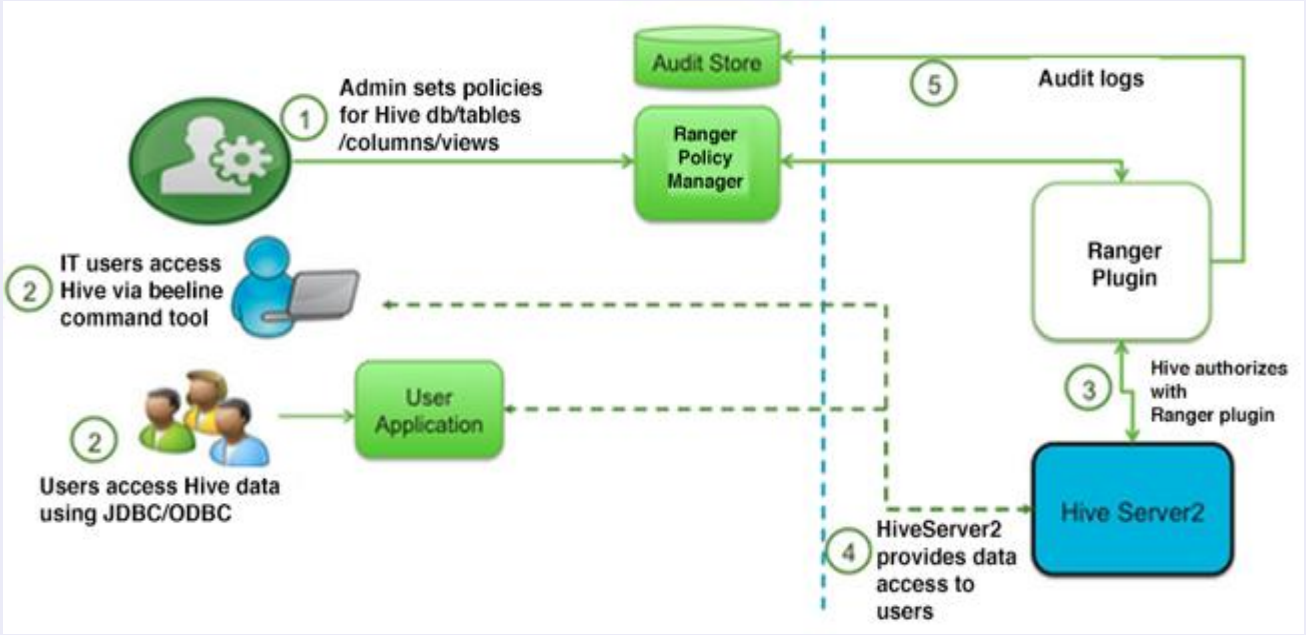

http://docs.hortonworks.com/HDPDocuments/HDP2/HDP-2.3.2/bk_Sys_Admin_Guides/content/ref-746ce51a-9bdc-4fef-85a6-69564089a8a6.1.html

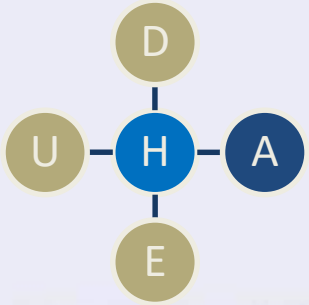

Apache Ranger overview

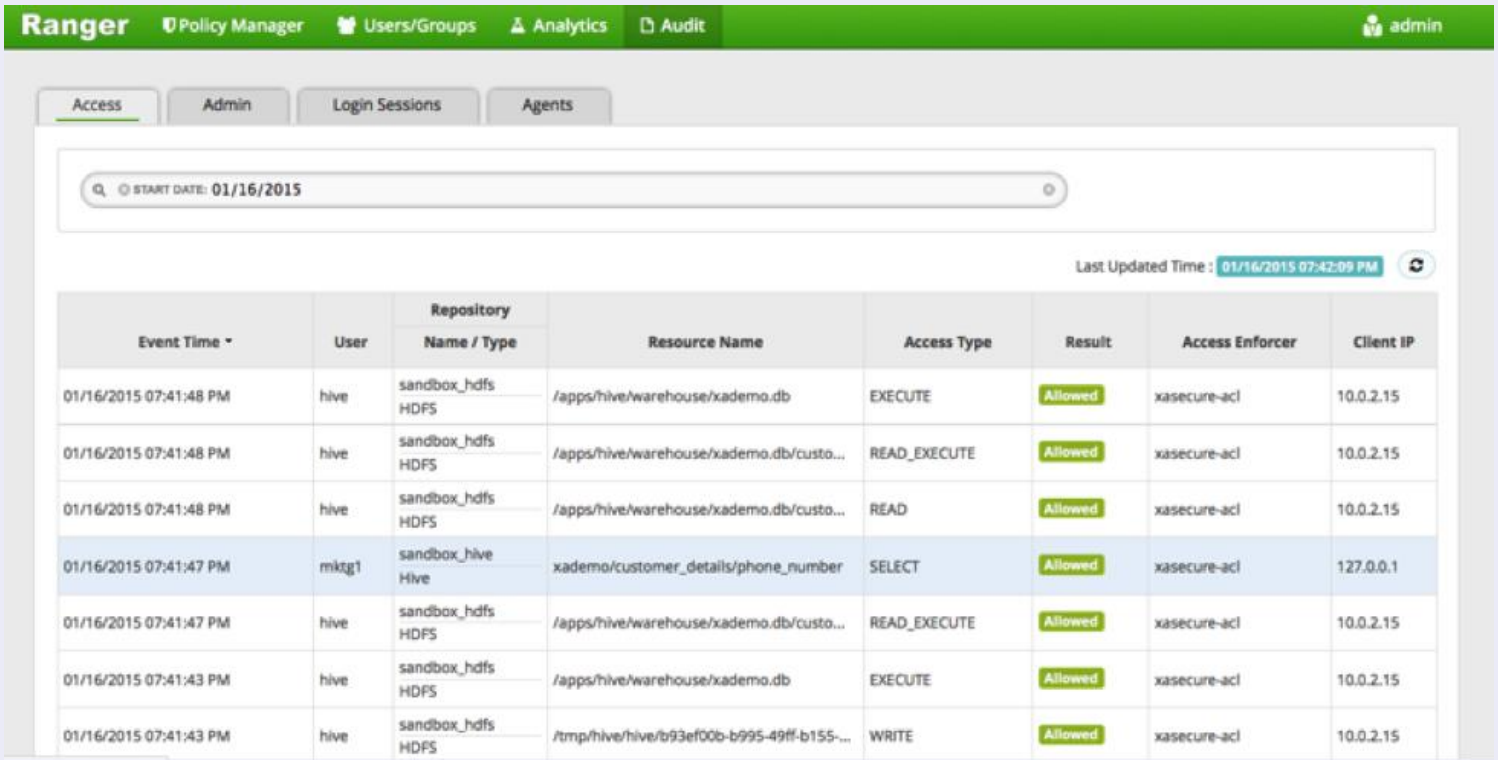

http://hortonworks.com/blog/best-practices-for-hive-authorization-using-apache-ranger-in-hdp-2-2/

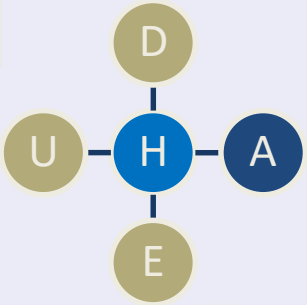

Apache Ranger

OWASP 中国
The Open Web Application Security Project

- Low hanging fruits: no HTTP hardening, SlowHTTP DoS
- Standard users can log into Ranger but have no permissions
- Interesting function level access control ->

Apache Ranger vulnerabilities

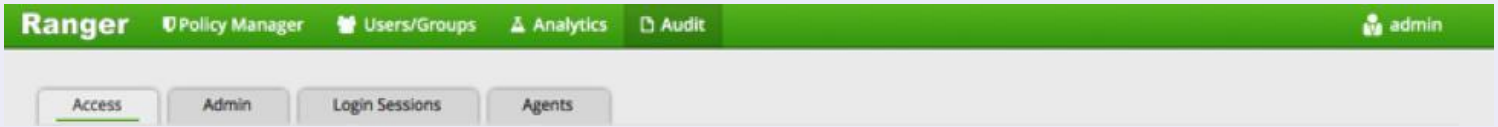


```
<a href="#!/policymanager" id="nav2"><i class="icon-shield"></i>{{tt 'h.policyManager'}} </a>
</li>
{{#isSystemAdmin .}}
<li>
<a href="#!/users/usertab" id="nav3"><i class="icon-group"></i> {{tt 'h.usersOrGroups'}} </a>
</li>
{{/isSystemAdmin}}
<li>
<a href="#!/reports/userAccess" id="nav7"><i class="icon-beaker"></i> {{tt 'h.analytics'}} </a>
</li>
{{#isSystemAdmin .}}
<li>
<a href="#!/reports/audit/bigData" id="nav8"><i class=" icon-file-alt"></i> {{tt 'h.audit'}} </a>
</li>
{{/isSystemAdmin}}
```

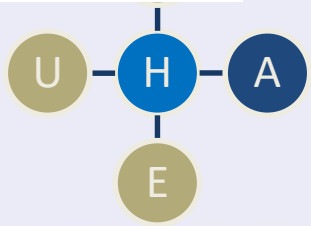

Missing function level access control

OWASP 中国
The Open Web Application Security Project

- Audit (X)
  - Big Data (X)
  - Admin (V)
  - Login Sessions (X)
    - Sessoin details (X)
    - Show actions (V)
- Users/Group (X)
  - Add new user (V)
  - List (X)
    - List (X)
    - Edit (V)
- Policies/Analytics (V)
  - List (V)
  - Edit (X)
    - Save changes (V)
    - Details (X)
  - Delete (X)

CVE-2015-0266

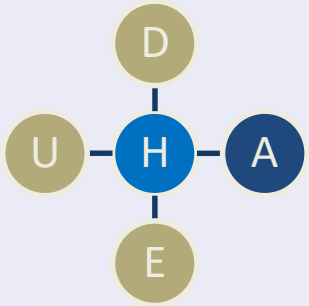

Apache Ranger attack scenario

OWASP 中国
The Open Web Application Security Project

- Target an old Hadoop installation (Apache Ranger 0.4 or XA-Secure v. 3.5.001)
- Hijack standard Hadoop account
- Log into Ranger (with low permissions)
- Use CVE-2015-0266 to escalate privileges
- Edit accounts, authorization rules, access policies

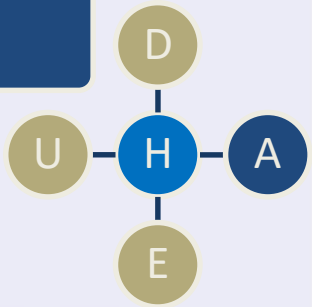

Apache Ranger vulnerabilities

OWASP 中国
The Open Web Application Security Project

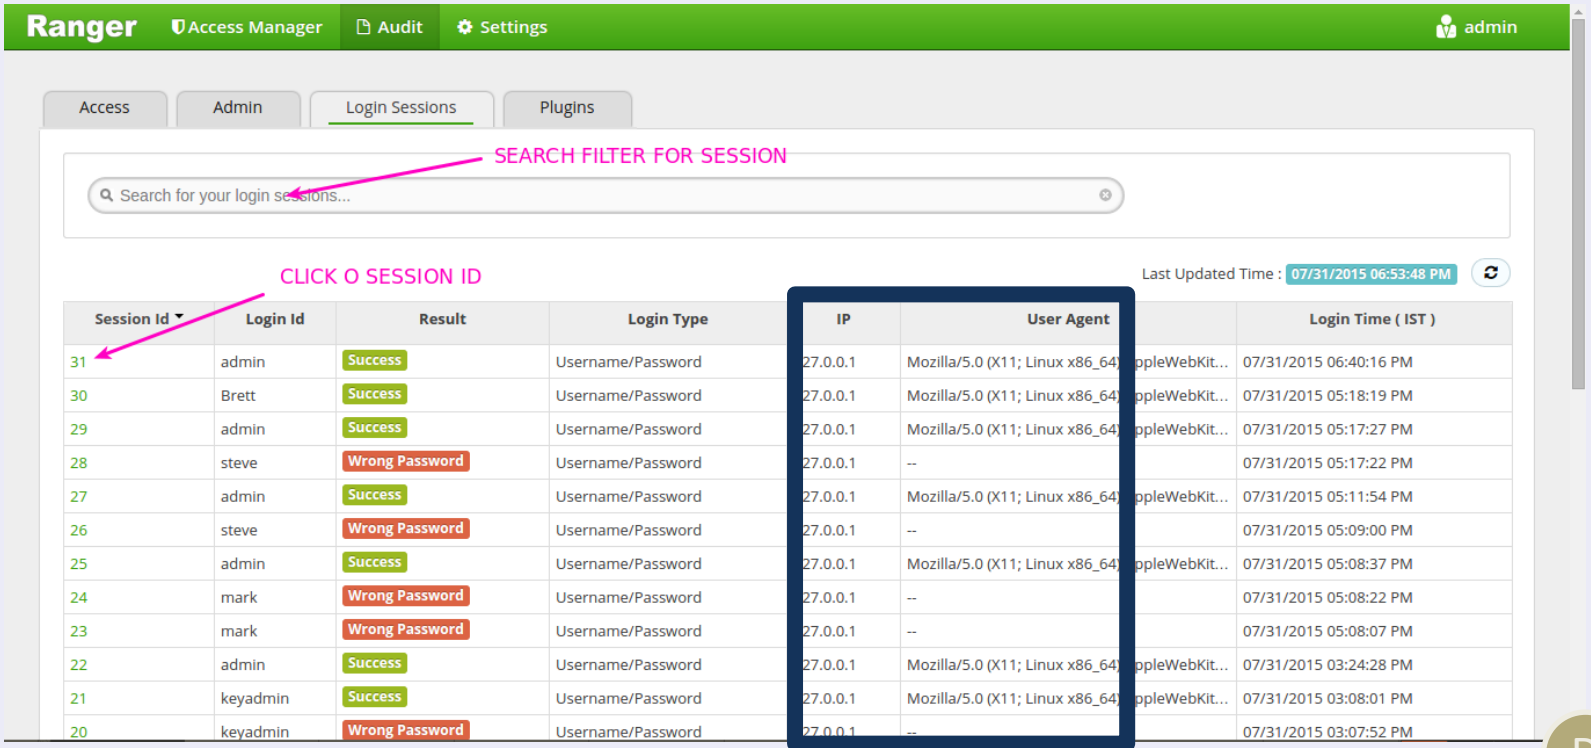

https://cwiki.apache.org/confluence/display/RANGER/Apache+Ranger+0.5+-+User+Guide

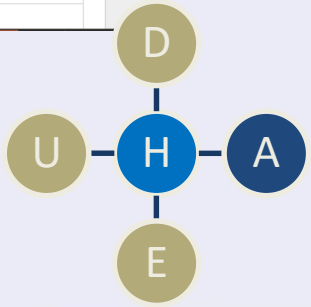

Apache Ranger XSS through UserAgent

OWASP 中国
The Open Web Application Security Project

- User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) <script>alert(1);</script>

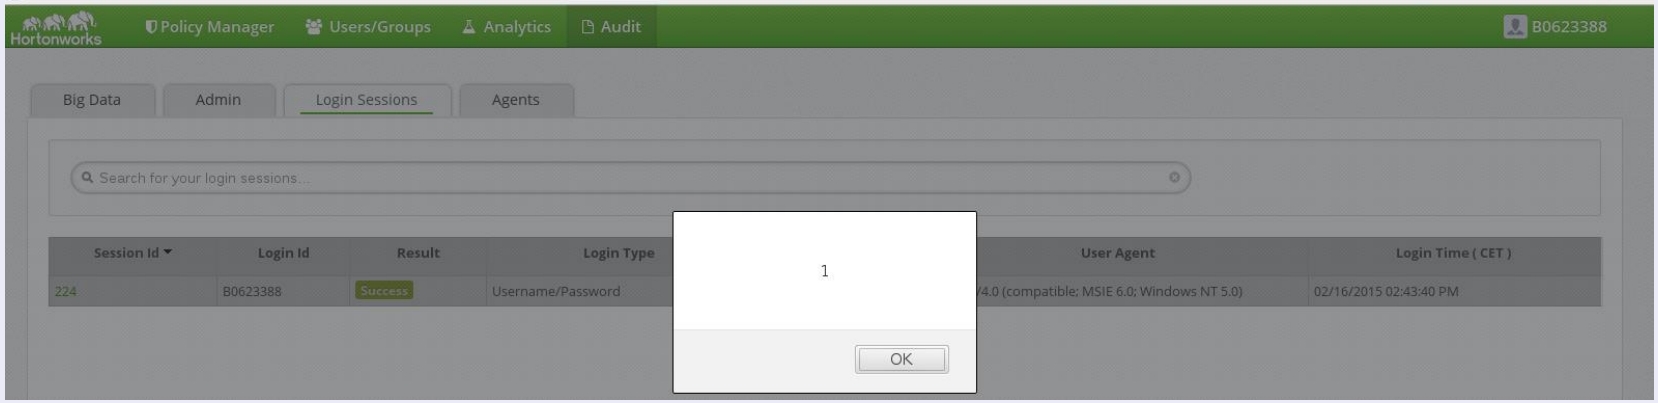

CVE-2015-0265

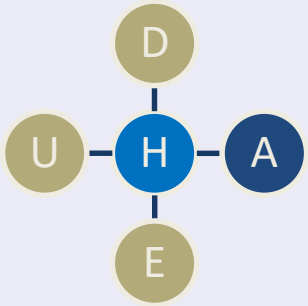

Apache Ranger attack scenario

OWASP 中国
The Open Web Application Security Project

Target an old Hadoop installation (Apache Ranger 0.4 or XA-Secure v. 3.5.001)

Network access to Apache Ranger is necessary (either from the internet or local network)

Log in with any user and password using XSS in UserAgent

You don´t need to escalate privileges, you´re already an admin (after admin opens session tab)

Deploy BEEF or whatsoever (CSRF script) to create users and change policies

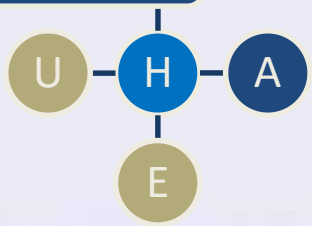

Apache Ranger patched

OWASP 中国
The Open Web Application Security Project

- Affected version: Apache Ranger v 0.4.0, XA Secure v. 3.5.001
- Both vulnerabilities patched in Ranger v 0.5.0
- For a while developers did a self-full-disclosure ->

RANGER-284 in public Jira now

OWASP 中国
The Open Web Application Security Project

Ranger / RANGER-284

Replace "Agents" with "Plugins" in Ranger Admin UI

Agile Board

Export

Details

Type:	Bug	Status:	RESOLVED
Priority:	Major	Resolution:	Fixed
Affects Version/s:	0.4.0	Fix Version/s:	0.5.0
Component/s:	None		
Labels:	None		

Description

Review all references to "Agent" in the UI templates and replace them with "Plugin". For Eg :
Page: Audit==>Agents:
Search text: "Search for your agents.."
Search fields: "Agent Id", "Agent IP"
Columns: "Agent Id", "Agent IP"

People

Assignee: Gautam Borad

Reporter: Gautam Borad

Votes: 0 Vote for this issue

Watchers: 1 Start watching this issue

Dates

Created:

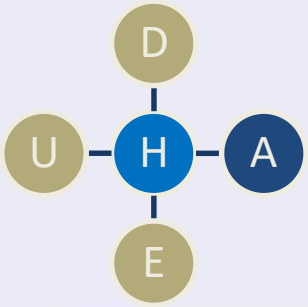

RANGER-284 shortly after vendor contact

OWASP 中国
The Open Web Application Security Project

```
Gautam Borad updated RANGER-284:
-------------------------------
    Attachment: RANGER-284-Escape-HTML-before-displaying-to-prevent-.patch

> Sanitize User Data to prevent XSS - Security Vulnerability
> ----------------------------------------------------------
>
>                 Key: RANGER-284
>                 URL: https://issues.apache.org/jira/browse/RANGER-284
>             Project: Ranger
>          Issue Type: Bug
>    Affects Versions: 0.4.0
>            Reporter: Gautam Borad
>            Assignee: Gautam Borad
>             Fix For: 0.5.0
>
>         Attachments: RANGER-284-Escape-HTML-before-displaying-to-prevent-.patch
>
>
> *Steps to reproduce*
> * Set user agent to something like this - "Mozilla/4.0 (compatible; MSIE 6.0; Windows
NT 5.0) <script>alert(1);</script>"
> * Try to login to policy admin with an incorrect username/password
> * Now login as admin user
> * Go to Audit tab --> Login Sessions
> * You will notice the failed logins displayed
> * Click on the failed login session id
> * Click Login sessions
> * You will notice a Javascript popup alert (entered in the user agent)
> *Expected Result*
> Unauthorized users should not be able to change the behavior of the application
> *Actual Result*
> Unauthorized users are able to put javascript code that can be executed in admin users
context
> *Fix*
> Sanitize the user input data and any data comes from user.
```

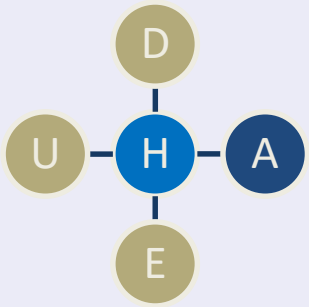

OWASP 中国
The Open Web Application Security Project

not in every environment

DISTRIBUTIONS SPECIFICS

Distribution specifics

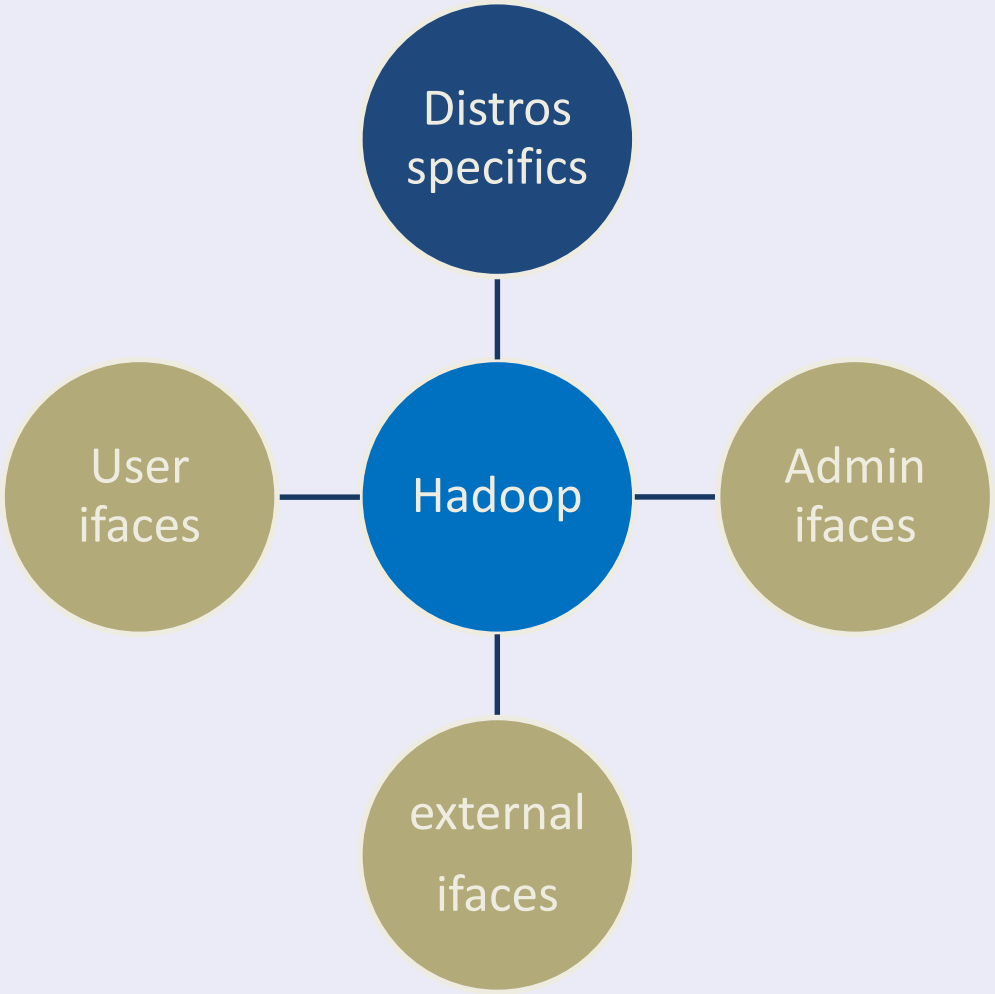

Distros

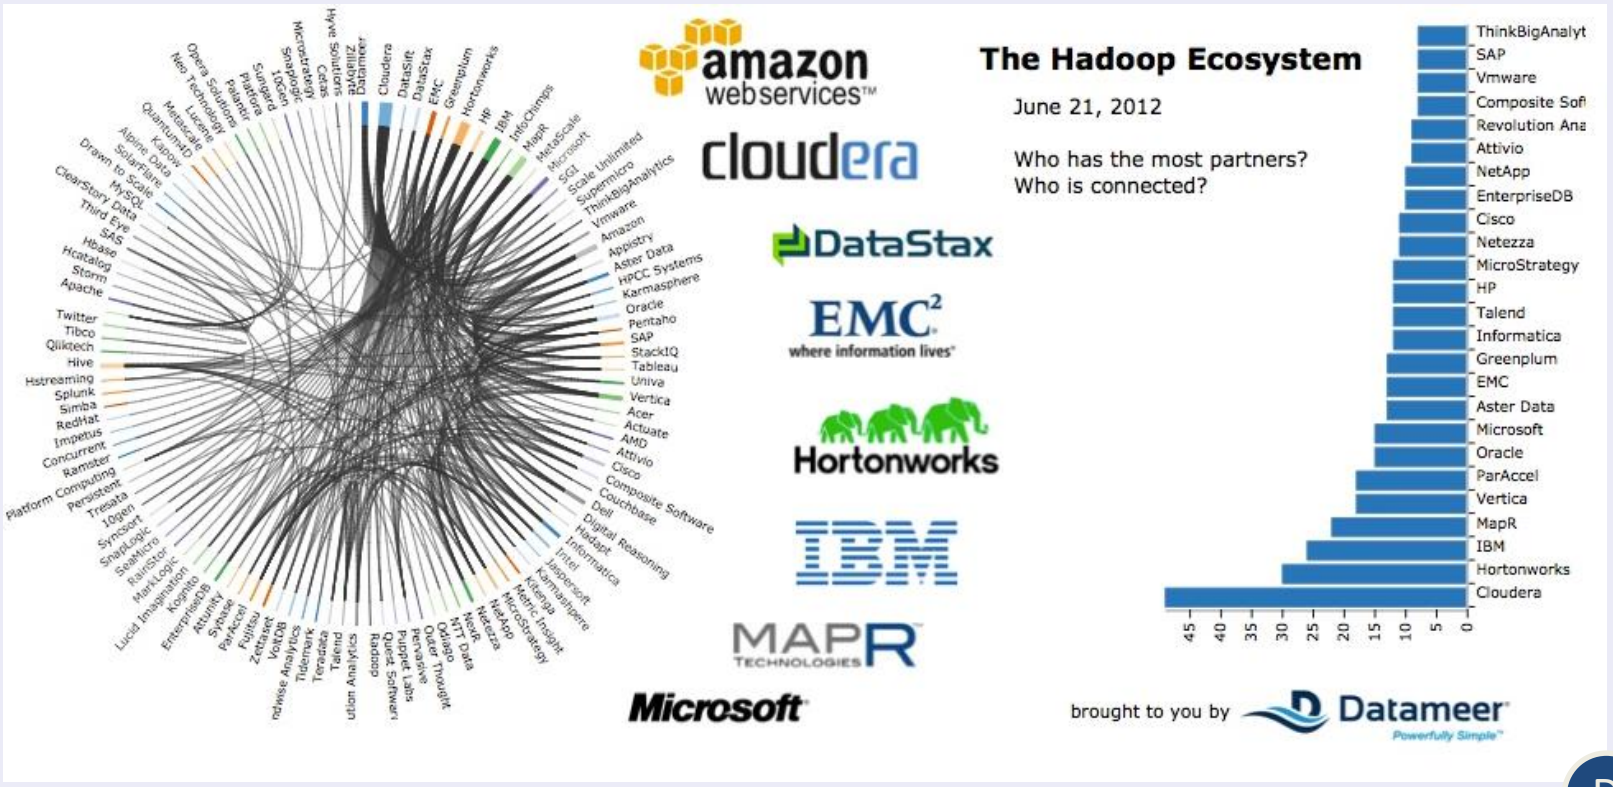

http://blog.cloudera.com/blog/2012/07/the-hadoop-ecosystem-visualized-in-datameer/

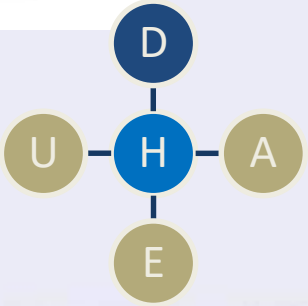

Basic distinction

OWASP 中国
The Open Web Application Security Project

cloud based	hosted locally

Distros

- How long does it take to create a new distro version?
- How many components are outdated at that time?
- How long does it take to deploy a new distro at a company?
- How many components are outdated at that time?

Most cases:

- MAJOR – ca. 1 year
- MINOR – ca. 3 months
- PATCH – ca. 1-2 months (differs much)

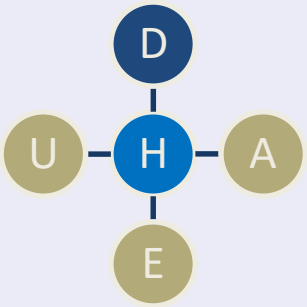

Hortonworks HDP components by version

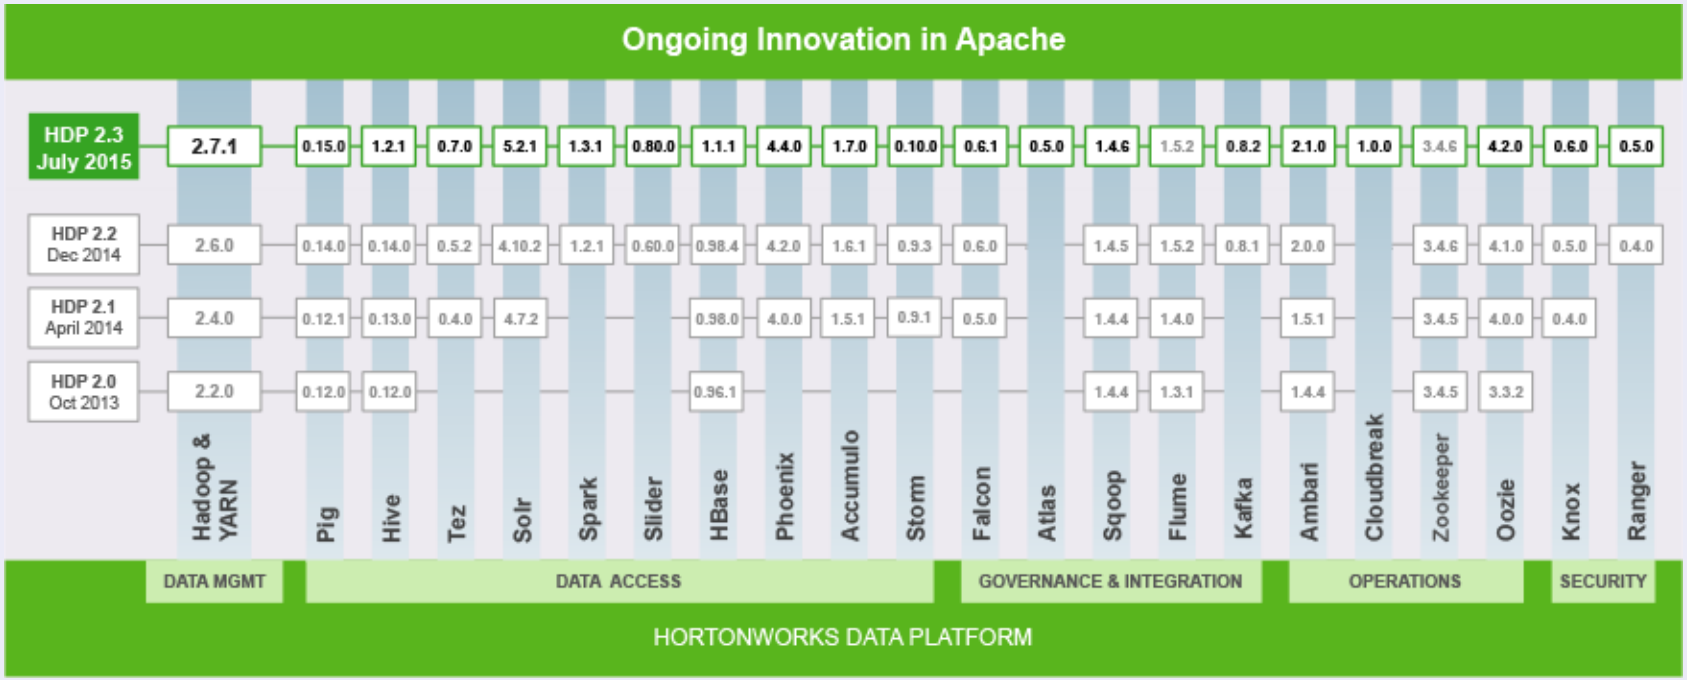

http://hortonworks.com/hdp/whats-new/

Distros

Old components with known issues

- Old OS components (java, php, ruby, etc.)
- Old OS components (e.g. old tomcat used by Oozie and HDFS)
- Old Hadoop components (e.g. old Hue, Ambari, Ranger)

Default passwords

Default configuration

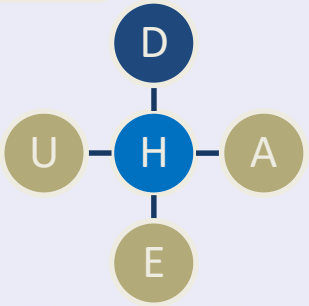

Vulnerability timeline

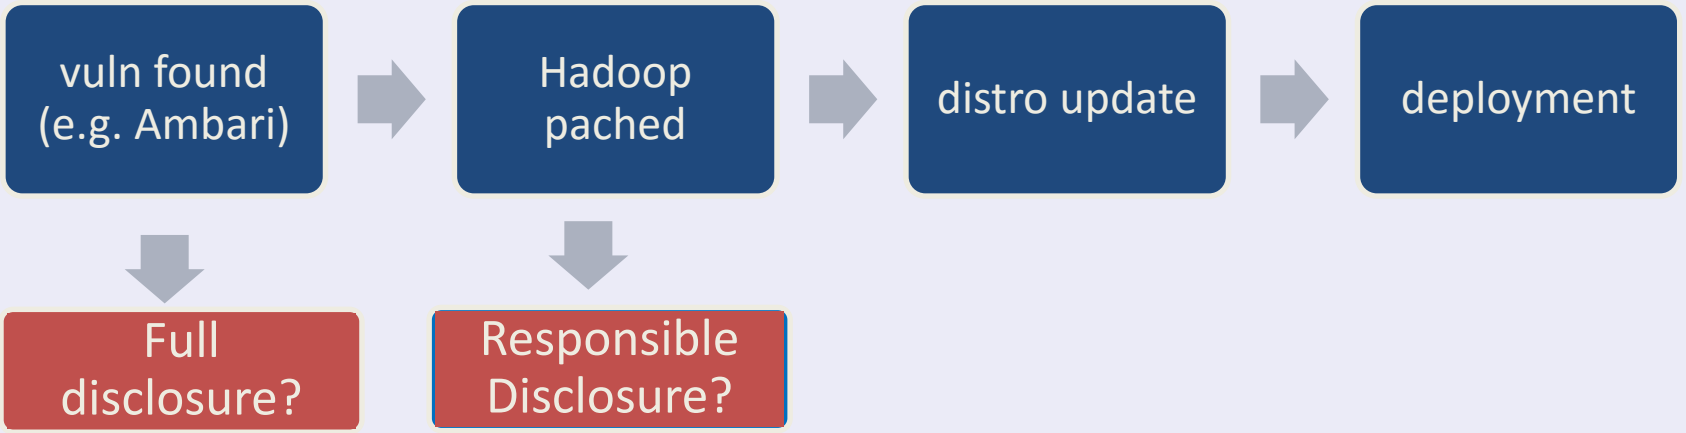

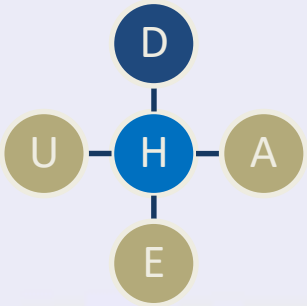

Distros

OWASP 中国
The Open Web Application Security Project

Old components with known issues

Default passwords

- SSH keys configured but default passwords still work
- Default mysql passwords, NO mysql passwords

Default configuration

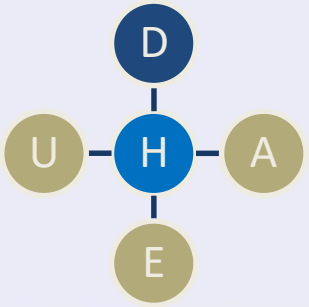

Distros

Old components with known issues

Default passwords

Default configuration

- No network level hardening
- No HTTP hardening (clickjacking, session mgmt, errors)
- Hue uses Django with DEBUG turned on by default
- „Hacking virtual appliances” by Jeremy Brown

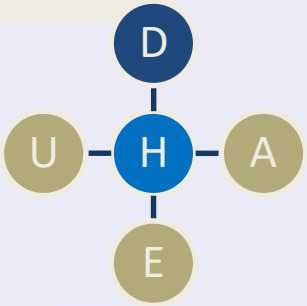

Default configurations sucks

OWASP 中国
The Open Web Application Security Project

X-Frame-Options:ALLOWALL

SESSION_COOKIE_DOMAIN	None
SESSION_COOKIE_NAME	'sessionid'
SESSION_COOKIE_PATH	'/;HttpOnly'
SESSION_COOKIE_SECURE	True
SESSION_ENGINE	'django.contrib.sessions.backends.db'
SESSION_EXPIRE_AT_BROWSER_CLOSE	False
SESSION_FILE_PATH	None
SESSION_SAVE_EVERY_REQUEST	False
SETTINGS_MODULE	'desktop.settings'
SHORT_DATETIME_FORMAT	'm/d/Y P'
SHORT_DATE_FORMAT	'm/d/Y'
SITE_ID	1
SKIP_SOUTH_TESTS	True
TEMPLATE_CONTEXT_PROCESSORS	('django.contrib.auth.context_processors.auth', 'django.core.context_processors.debug', 'django.core.context_processors.media', 'django.contrib.messages.context_processors.mess
TEMPLATE_DEBUG	False
TEMPLATE_DIRS	('/usr/lib/hue/desktop/core/templates',)
TEMPLATE_LOADERS	('django.template.loaders.filesystem.load_template_source', 'desktop.lib.template_loader
TEMPLATE_STRING_IF_INVALID	''
TEST_DATABASE_CHARSET	None
TEST_DATABASE_COLLATION	None
TEST_DATABASE_NAME	None
TEST_RUNNER	'django.test.simple.DjangoTestSuiteRunner'
THOUSAND_SEPARATOR	','
TIME_FORMAT	'P'
TIME_INPUT_FORMATS	('%H:%M:%S', '%H:%M')
TIME_ZONE	'America/Los_Angeles'
TRANSACTIONS_MANAGED	False
URL_VALIDATOR_USER_AGENT	'Django/1.2.3 (http://www.djangoproject.com)'
USE_ETAGS	False
USE_I18N	True
USE_L10N	True
USE_THOUSAND_SEPARATOR	False
X_FRAME_OPTIONS	'ALLOWALL'
YEAR_MONTH_FORMAT	'F Y'

You're seeing this error because you have DEBUG = True in your Django settings file. Change that to False, and Django will display a standard 500 page.

debug 7 of 7

Page not found (404)

Request Method: GET
Request URL: https://lxhulo001:8000/x

Using the URLconf defined in desktop.urls, Django tried these URL p

1. ^about/static/(?P<path>.*)$
2. ^beeswax/static/(?P<path>.*)$
3. ^filebrowser/static/(?P<path>.*)$
4. ^hcatalog/static/(?P<path>.*)$
5. ^help/static/(?P<path>.*)$
6. ^jobbrowser/static/(?P<path>.*)$
7. ^jobsub/static/(?P<path>.*)$
8. ^oozie/static/(?P<path>.*)$
9. ^pig/static/(?P<path>.*)$
10. ^shell/static/(?P<path>.*)$
11. ^useradmin/static/(?P<path>.*)$
12. ^static/(?P<path>.*)$
13. ^(?P<path>favicon.ico)$
14. ^accounts/login/$
15. ^accounts/logout/$
16. ^logs$
17. ^dump_config$
18. ^download_logs$
19. ^bootstrap.js$
20. ^profile$
21. ^prefs/(?P<key>\w+)?$
22. ^status_bar/?$
23. ^admin/
24. ^debug/threads$
25. ^debug/who_am_i$
26. ^debug/check_config$
27. ^debug/check_config_ajax$
28. ^log_frontend_event$
29. ^jasmine
30. ^$
31. ^about/
32. ^beeswax/
33. ^filebrowser/
34. ^hcatalog/
35. ^help/
36. ^jobbrowser/
37. ^jobsub/
38. ^oozie/
39. ^pig/

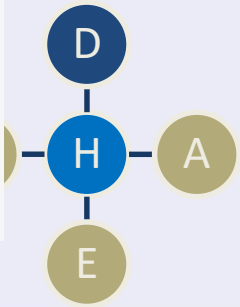

OWASP 中国
The Open Web Application Security Project

For clients or whatsoever

EXTERNAL INTERFACES

External interfaces

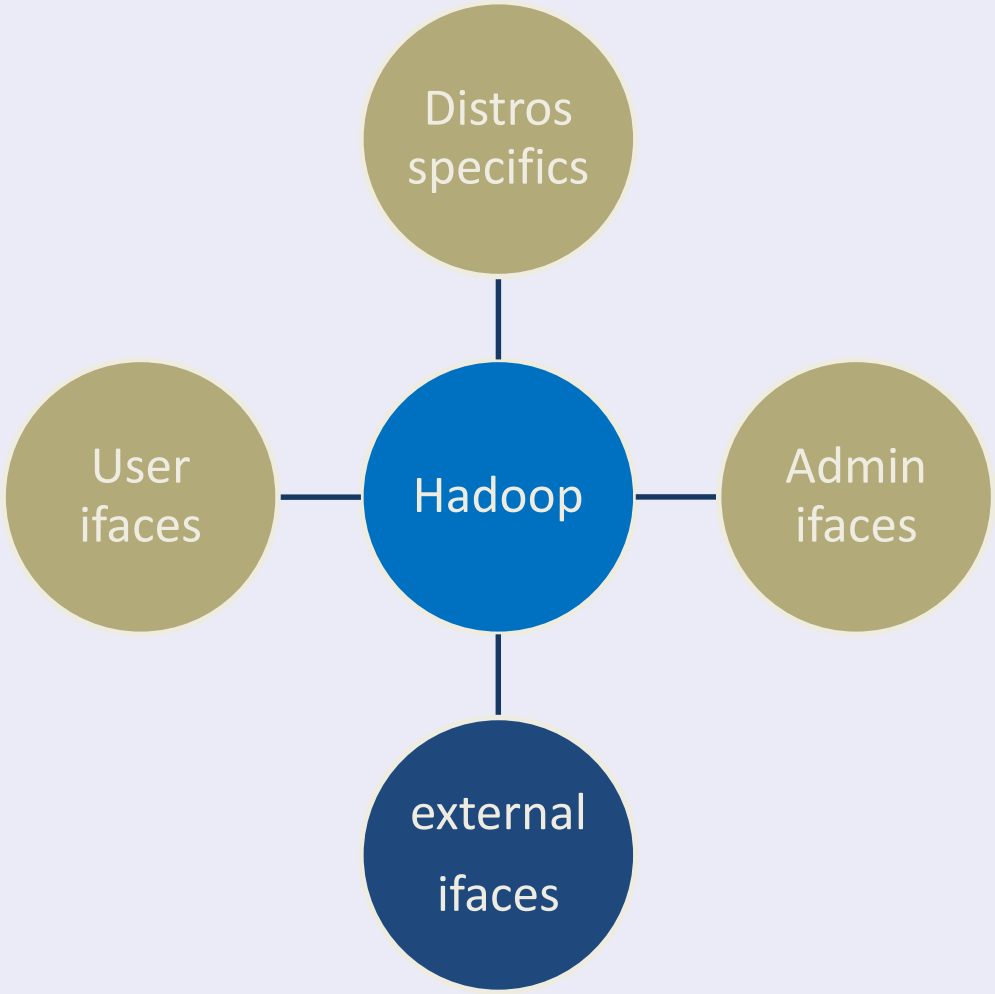

External

- More than 25 internal Apache apps/modules
- Vendor/distro specific apps/interfaces
- Popular monitoring: Ganglia, Splunk
- Auth providers: LDAP, Kerberos, OAuth
- Many apps, many targets

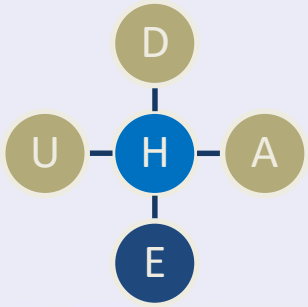

Hadoop

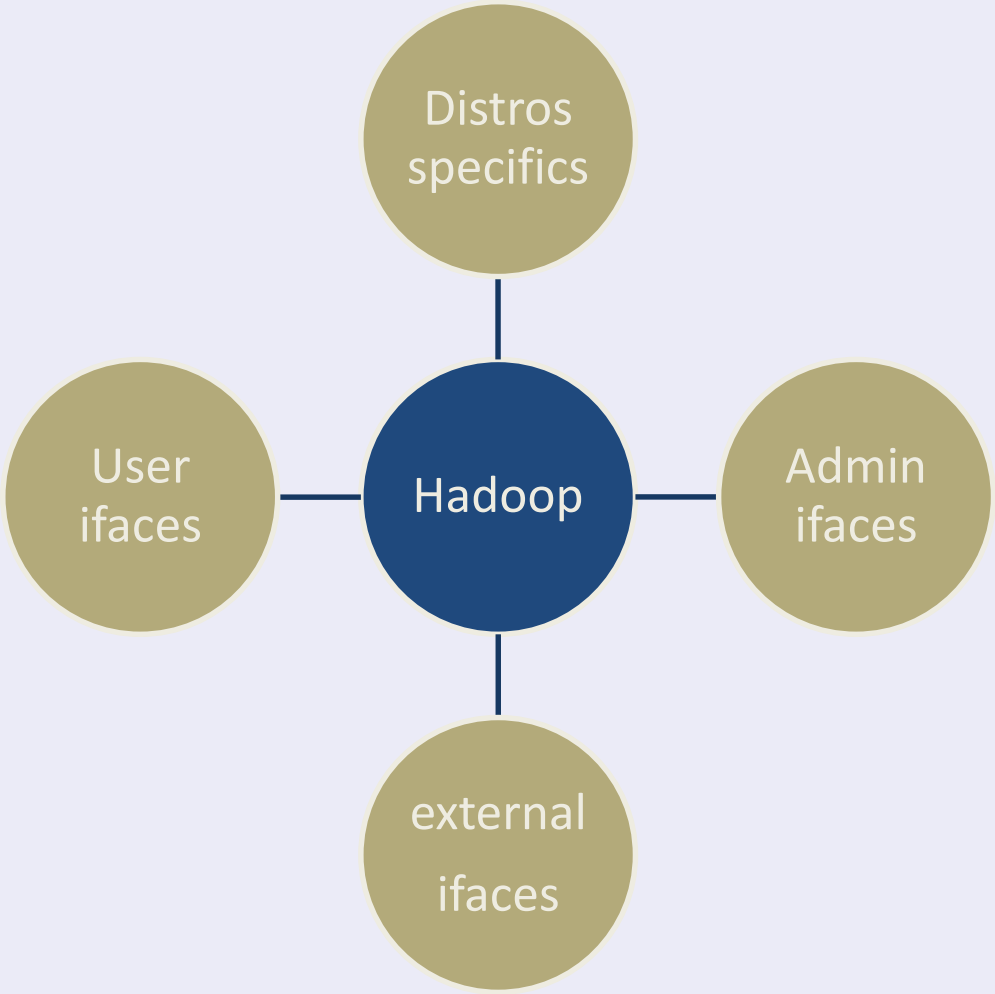
Distros specifics
User ifaces
Hadoop
Admin ifaces
external ifaces

OWASP 中国
The Open Web Application Security Project

ways to protect your big data environment

SUMMARY

Ways to protect your Hadoop environment

OWASP 中国
The Open Web Application Security Project

Excessive network access

- Keep it super tight!

Excessive user pesmissions

Typical web vulnerabilities

Obsolete software

Distros dependent vulnerabilities

External system connections

Ways to protect your Hadoop environment

OWASP 中国
The Open Web Application Security Project

Excessive network access

Excessive user permissions

- Map business roles to permissions

Typical web vulnerabilities

Obsolete software

Distros dependent vulnerabilities

External system connections

Ways to protect your Hadoop environment

OWASP 中国
The Open Web Application Security Project

Excessive network access

Excessive user permissions

Typical web vulnerabilities

- Pentest it! Introduce application independent security countermeasures

Obsolete software

Distros dependent vulnerabilities

External system connections

Ways to protect your Hadoop environment

OWASP 中国
The Open Web Application Security Project

Excessive network access

Excessive user permissions

Typical web vulnerabilities

Obsolete software

- Make a list of all components. Monitor bugtracks and CVEs.

Distros dependent vulnerabilities

External system connections

Ways to protect your Hadoop environment

OWASP 中国
The Open Web Application Security Project

Excessive network access

Excessive user permissions

Typical web vulnerabilities

Obsolete software

Distros dependent vulnerabilities

- A pentest after integration is a must. Demand security from software suppliers.

External system connections

Ways to protect your Hadoop environment

OWASP 中国
The Open Web Application Security Project

- Excessive network access
- Excessive user permissions
- Typical web vulnerabilities
- Obsolete software
- Distros dependent vulnerabilities
- External system connections

- Make a list of all external system connections. Do a threat modeling and pentest corresponding systems.

Current status

OWASP 中国
The Open Web Application Security Project

Apache » Hadoop : Vulnerability Statistics

Apache » Hive : Vulnerability Statistics

Apache » Hbase : Vulnerability Statistics

Apache » Ambari : Vulnerability Statistics

Apache » Ranger : Vulnerability Statistics

Apache » Cassandra : Vulnerability Statistics

Vulnerabilities (1) CVSS Scores Report Browse all versions Possible matches for this product Related Metasploit Modules

Related OVAL Definitions : Vulnerabilities (0) Patches (0) Inventory Definitions (0) Compliance Definitions (0)

Vulnerability Feeds & Widgets

Vulnerability Trends Over Time

Year	# of Vulnerabilities	DoS	Code Execution	Overflow	Memory Corruption	Sql Injection	XSS	Directory Traversal	Http Response Splitting	Bypass something	Gain Information	Gain Privileges	CSRF	File Inclusion	# of exploits
2015	1		1												
Total	1		1												
% Of All		0.0	100.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	

- Spark™: A fast and general compute engine for Hadoop data. Spark provides a simple and expressive programming model that s computation.
- Tez™: A generalized data-flow programming framework, built on Hadoop YARN, which provides a powerful and flexible engine to e being adopted by Hive™, Pig™ and other frameworks in the Hadoop ecosystem, and also by other commercial software (e.g. ETL
- ZooKeeper™: A high-performance coordination service for distributed applications.

OWASP 中国
The Open Web Application Security Project

谢谢

THANK YOU!

Jakub Kaluzny