

Taming Vulnerability Data

Vulnerability scanners are indispensable both for vulnerability assessments and penetration tests. One of the first things a tester does when faced with a network is fire up a network scanner or even several different ones. And, having scanned a large network with a lot of open ports and services he/she will end up with something like this: see Figure 1.

This is however only the first step. The next steps would usually be:

- Get a general idea of what is there. Are there any problems that look exploitable? What are common problems? What hosts are affected?
- Weed out the findings that are obviously bogus or unimportant.
- Attempt exploitation to verify the findings. In case of a penetration test this step will include exploiting, establishing pivot points, escalating privileges, gathering further information, and so on. In any case it will involve exploiting vulnerabilities and recording the results of the exploitation.
- Gather additional information if something is missing

Host	Findings
.100	Low Severity problem(s) found
.101	High Severity problem(s) found
.97	Medium Severity problem(s) found
.98	Low Severity problem(s) found
.99	High Severity problem(s) found
.193	Low Severity problem(s) found
.194	Low Severity problem(s) found
.195	Low Severity problem(s) found
.196	Medium Severity problem(s) found
.197	High Severity problem(s) found
.198	Low Severity problem(s) found
.1	Low Severity problem(s) found
.10	Low Severity problem(s) found
.101	Low Severity problem(s) found
.102	Low Severity problem(s) found
.105	Low Severity problem(s) found
.106	Low Severity problem(s) found
109	Low Severity problem(s) found

Figure 1. Nessus report. 404 hosts, 960 open ports, 3286 findings

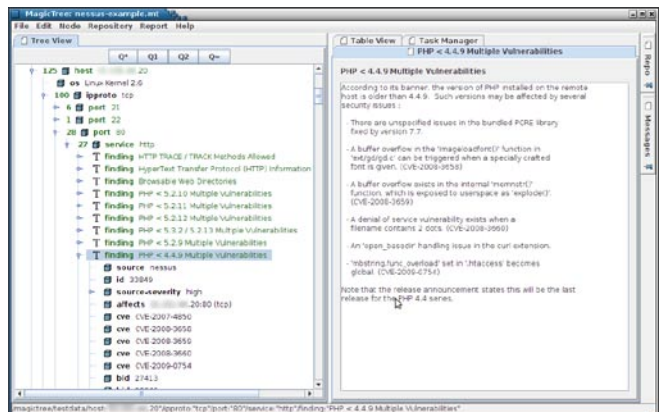


Figure 2. MagicTree shows a host, ports, service and finding in Tree View

Severity	Affected host	Affected port
High	192.168.1.1	21
High	192.168.1.1	22
High	192.168.1.1	23
High	192.168.1.1	24
High	192.168.1.1	25
High	192.168.1.1	26
High	192.168.1.1	27
High	192.168.1.1	28
High	192.168.1.1	29
High	192.168.1.1	30
High	192.168.1.1	31
High	192.168.1.1	32
High	192.168.1.1	33
High	192.168.1.1	34
High	192.168.1.1	35
High	192.168.1.1	36
High	192.168.1.1	37
High	192.168.1.1	38
High	192.168.1.1	39
High	192.168.1.1	40
High	192.168.1.1	41
High	192.168.1.1	42
High	192.168.1.1	43
High	192.168.1.1	44
High	192.168.1.1	45
High	192.168.1.1	46
High	192.168.1.1	47
High	192.168.1.1	48
High	192.168.1.1	49
High	192.168.1.1	50
High	192.168.1.1	51
High	192.168.1.1	52
High	192.168.1.1	53
High	192.168.1.1	54
High	192.168.1.1	55
High	192.168.1.1	56
High	192.168.1.1	57
High	192.168.1.1	58
High	192.168.1.1	59
High	192.168.1.1	60
High	192.168.1.1	61
High	192.168.1.1	62
High	192.168.1.1	63
High	192.168.1.1	64
High	192.168.1.1	65
High	192.168.1.1	66
High	192.168.1.1	67
High	192.168.1.1	68
High	192.168.1.1	69
High	192.168.1.1	70
High	192.168.1.1	71
High	192.168.1.1	72
High	192.168.1.1	73
High	192.168.1.1	74
High	192.168.1.1	75
High	192.168.1.1	76
High	192.168.1.1	77
High	192.168.1.1	78
High	192.168.1.1	79
High	192.168.1.1	80
High	192.168.1.1	81
High	192.168.1.1	82
High	192.168.1.1	83
High	192.168.1.1	84
High	192.168.1.1	85
High	192.168.1.1	86
High	192.168.1.1	87
High	192.168.1.1	88
High	192.168.1.1	89
High	192.168.1.1	90
High	192.168.1.1	91
High	192.168.1.1	92
High	192.168.1.1	93
High	192.168.1.1	94
High	192.168.1.1	95
High	192.168.1.1	96
High	192.168.1.1	97
High	192.168.1.1	98
High	192.168.1.1	99
High	192.168.1.1	100

Figure 3. MagicTree TableView lists of findings by severity and type

- Produce a human readable report including only the relevant findings. The *human readable* part is what makes it difficult. By that we mean a report that a human, such as a security officer or a system administrator will actually read it in full or easily find the parts that are relevant to him/her without missing anything important. A thousand page report where findings are listed by host with tens of findings for each of a hundred of hosts does not qualify as such.

Unfortunately, vulnerability scanners, such as Nessus or OpenVAS, stop at producing the report. What you do with this data and how you do it is up to you.

Having faced these problems regularly, we came up with a solution. Enter MagicTree.

Making Sense of Data

In vulnerability assessments we deal with various kinds of objects such hosts, ports, services, applications, and vulnerabilities. Hosts have IP addresses, DNS names, operating systems, and various other attributes. They also have TCP/UDP ports, that may be open, closed or filtered. Ports have services or protocols, and also software listening on those ports. And all those might have vulnerabilities.

You will notice that these objects are naturally organized in a tree-like structure. This is the structure that MagicTree uses to represent the data: see Figure 2.

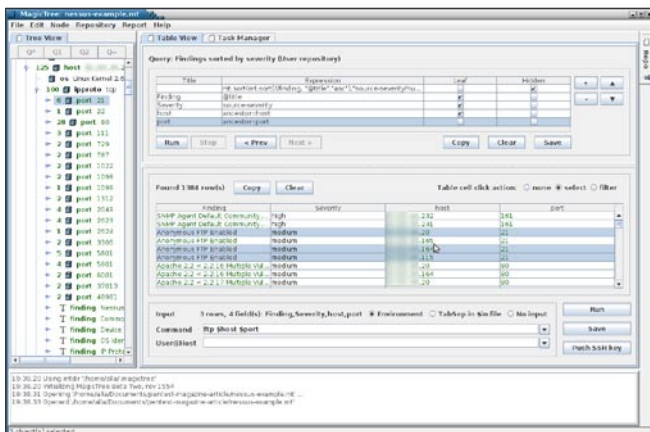


Figure 4. Selecting systems that have anonymous FTP

The data can come from a variety of sources, including vulnerability scanners, flat files and even manual entry. In this case we are concerned with vulnerability scanners. Figure 2 shows the data imported from an XML file generated by Nessus.

Another way to look at the vulnerability data is by type of the vulnerability. We want to see what kinds of vulnerabilities were found (ordered by severity) and which hosts are affected by each type of vulnerability. This will give us the general idea of how bad (or how good, depending on your point of view) things are, and what can be exploited. Having this kind of view we can proceed to exploiting individual vulnerabilities. MagicTree lets you do just that.

MagicTree allows querying data in the tree using XPath expressions. The data shown in Figure 3 is a result of such a query. A number of useful queries, such as this one, a pre-defined and stored in the query repository. Users can save queries they define to the repository and reuse them.

Using the Data

The tester's next step will usually be exploiting or verifying the discovered vulnerabilities. As an example let's consider anonymous FTP servers. In our scan, several FTP servers that allow anonymous access were discovered. We want to investigate those servers to see what files are accessible. In addition to that we want to run a password guessing attack on those servers to try to get authenticated FTP access. Let's also assume that the test scope was set, we were given a list of hosts that are critical and fragile, so no intrusive tests should be performed on them. One of our anonymous FTP servers is in the list, so we must exclude it from password brute forcing attack.

In the list of findings we can select all *Anonymous FTP Server* findings. The table contains the sever IP addresses and port numbers. One of the servers is

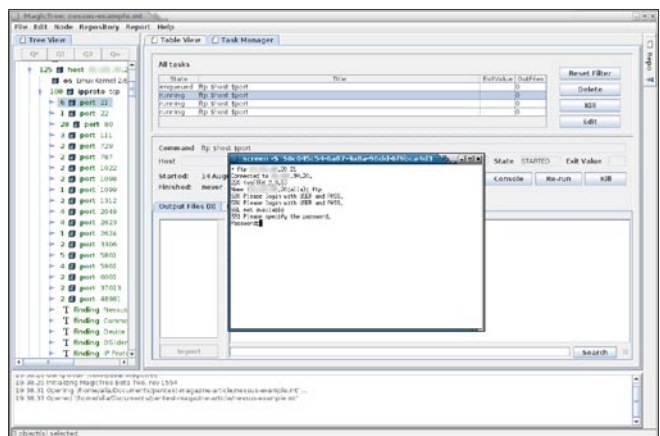


Figure 4. Running ftp for selected hosts and ports and opening FTP command console

out of bounds for us, so we control click on the row corresponding to that server and de-select it.

First we are going to connect with FTP to each of the servers and have a look around. We want to avoid having to copy and paste the IP addresses and port numbers, and then having to copy and paste the output into the report.

In MagicTree, under the table of results, there is an input field for the command line. We put in `ftp $host $port`, select *Environment* input mode and click Run. This will start one FTP command for each host and port in the list. Next we can open console for each running FTP session and look around.

When we terminate the FTP command, the console window closes, but the command output is saved. The output of all commands executed from MagicTree can be included in the generated report. Of course, commands that produced no useful output or were executed by mistake can be removed.

In a similar way we can run a password bruteforcer, such as hydra or medusa, feeding it the list of IP addresses, or running one instance per host. Again, all the output is retained by MagicTree and can be used in reports.

Advanced Commands Execution

The example above has shown some features related to the execution of external commands. We saw MagicTree launching a command-line FTP client, and passing target specification in environment variables, named after columns in the table view. The commands were executed locally, in parallel in small batches. Now we will consider slightly more complicated example. Suppose you have just scanned a bulk of class C networks using nmap, but only for hosts running web servers. We will want it to feed large bulk of data to nmap and execute the scanner on a remote host.

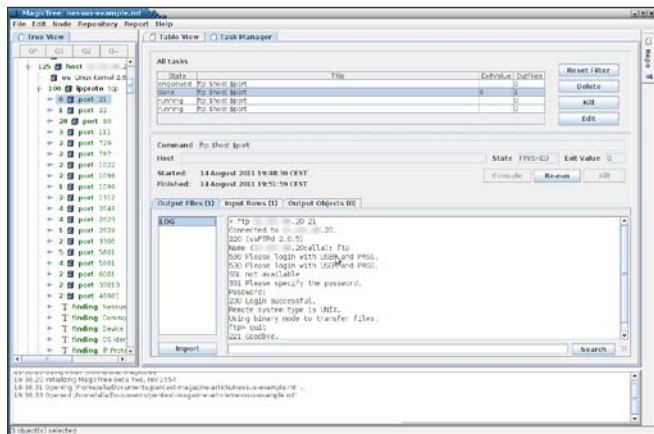


Figure 6. Browsing the FTP directories – screen log




Figure 7. Automatic query using Q2

Preparing Input Data

First we need to mine our tree for IP addresses of hosts having HTTP services open. One way to do it is to manually find an one TCP/80 port and then running a query to find similar nodes in the tree. This is done by searching for 80 in the tree and click on Q2 button above the tree. This button activates Query Wizard which builds a query to select all nodes in the tree having the same values in the two last nodes in the path (`port=80 and ipproto=tcp`).

Another option is to find all ports where HTTP service has been detected, plain or wrapped in SSL. The query for doing the latter is provided in the repository, just click on *Repo* tab, and find *HTTP and HTTPS servers* query under *ports* tag. Double click on the query and it will be automatically loaded and executed into the table view.

Launching the Command

Command editor is located under the table containing the query results. In the command editor we switch input mode to *Tabsep in \$in file*, and type the command:

```
sudo nmap -iL $in -sS -sV -n -vvv -p1-65535 -oX $out.xml".
```

When we launch the command, MagicTree will create a temporary file with the list of IP addresses and pass it to nmap via `$in` parameter. The command will run in background, detached from MagicTree, but you will have access to the interactive console which comes handy if you need enter password for *sudo*.

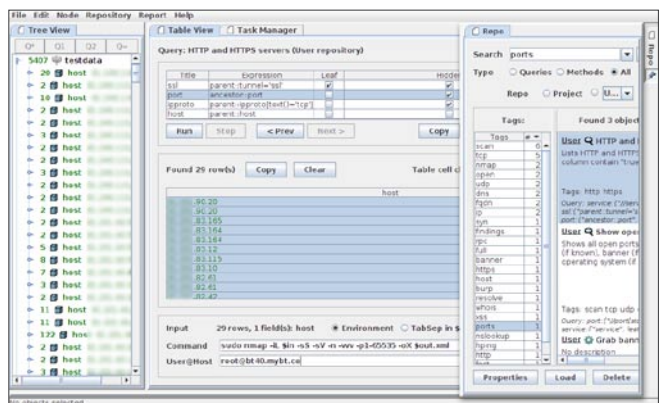


Figure 8. Ready to run the full port scan of web servers

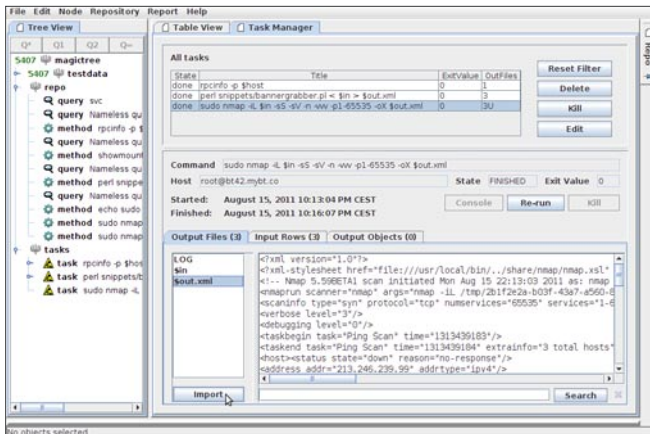


Figure 9. Output files produced by command can be imported back into the tree

After a command is launched, you can save the project, quit MagicTree and go home. Later you can open the project again and MagicTree will automatically check the status of commands it has launched and pick up its output files if the work is done.

About Output Files

MagicTree will retain the data the command outputs to the console. Additionally, it provides special parameter `$out`. It can be used in output file name parameters. In the above example we have passed `-oX $out.xml` to nmap. This will let MagicTree pick up the file produced by Nmap and store it. You could have written `-oA $out` instead and have all kind of nmap output stored in the tree, not only XML.

When the command finishes, you can browse the list of output files and their content. For XML files, you can click on Import button to import it back into the tree.

Taking import of nmap XML as an example, newly discovered ports will appear in the tree next to the ones previously discovered by Nessus.

Some tools such as DnsEnum and WhatWeb natively produce XML files in MagicTree format. As of time of writing this article, MagicTree supported XML files generated by Nessus, Nmap, Nikto, Acunetix, Burp, W3AF, Qualys, and OpenVAS 3. Users can develop XSLT transforms to support other tools that generate XML. Output of commands incapable of producing XML by themselves, can be imported into MagicTree by creating a wrapper.



Figure 10. Configuring remote BackTrack server to run commands from MagicTree




Figure 11. Marking out findings that we don't want in the final report

Remote Command Execution

Most likely you have already noticed `User@host` input fields beneath `Command`. It is related to another feature we would like to mention, remote command execution. Everything we have already said about running command from MagicTree – parameterization via environment variables or `$in`, output files pickup, XML file import, access to interactive console, detached execution – will work if a command is executed remotely. MagicTree uses no agents or third-party software installed on the remote host, it just needs to be able to login over SSH and find common *nix utilities.

To setup remote access, click on `Push SSH Key` button. A script running in newly open terminal window will generate a private SSH key (stored under `~/ .magictree`) and push it into authorized keys of the specified remote SSH server. How it looks like when I configure MagicTree to work with remote BackTrack 5 VPS, from HackingMachines: (Figure 10). Pushing has succeeded, so we can just enter `root@bt42.mybt.co` in `User@host` and have the command executed remotely.

Preparing the Report

Another challenge in processing the vulnerability scanner results is weeding out the findings that are useless, unimportant or plain wrong. Those will include findings that have no security impact and are no use to the client (stuff like *Traceroute information* or *Nessus scan information*) and false positives. Again MagicTree can help.

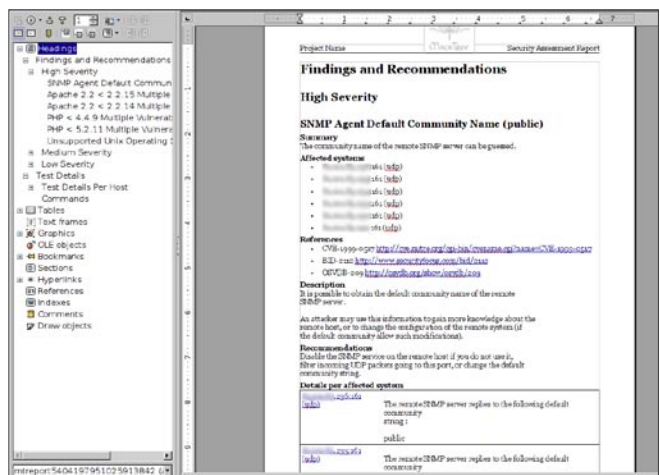


Figure 12. Generated report

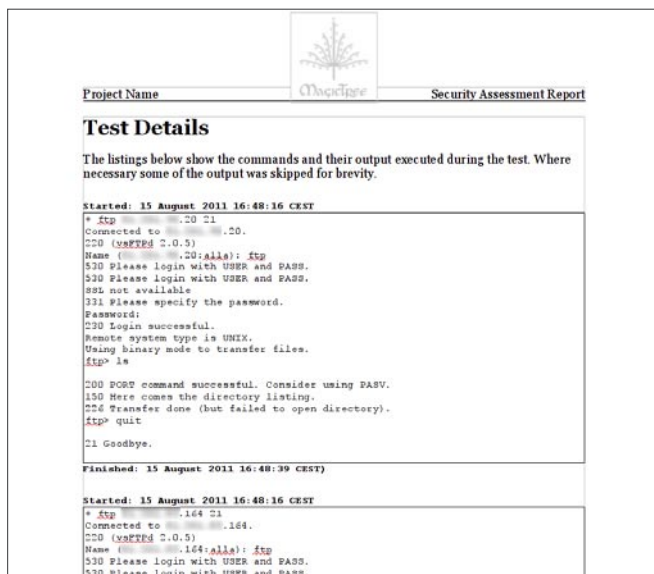


Figure 13. Command listings in the report

While deciding what is important and worth reporting and what isn't should be left to the human tester, the process of marking the findings as important or not and keeping only the important one in the final report, can and should be made easier.

Let's see how this can be done with MagicTree. We'll go back to the list of findings. Starting from the bottom, where all the low risk problems are, we will select findings we consider irrelevant and change their status to *ignore*.

The color of the table cells reflects the status. If we don't want to see findings marked as ignored in the list at all, all we need to do is change the query, adding a condition `[@status!='ignored']`, and re-run it.

Report Generation

After all this hard work sorting out and verifying the findings we want a report to show to the customer. The problem with the reports that vulnerability scanners, such as Nessus, generate is that they are difficult or impossible to edit and don't allow for much customization. MagicTree solves this problem, firstly by producing reports in Microsoft Office or OpenOffice Writer formats, so they are fully editable, and secondly providing fully customizable templates (also in Word and OpenOffice format), that can be modified to your liking. The templates use the same XPath expression syntax as the queries, so any data in the tree can be included in the report.

Let's start with one of the standard report templates that comes with MagicTree. Selecting *Report->Generate Report* from the menu, choosing *summary-of-findings-cross-referenced.odt* template and clicking on the Generate Report button produces the report: (Figure 12).

The first section – Findings and Recommendations – lists the findings grouping them by type, so *Anonymous FTP* is one finding containing the list of affected hosts and Nessus plugin output for each host. The second section – Test Details – contains per host data including the list of findings per host cross-referenced to the first section, and the listing of executed commands. Our FTP commands and brute-forcing sessions end up in this last section: (Figure 13).

In many cases the output of the vulnerability scanner is used during the test but does not make it into the final report, that is custom-written. Still there is often a need to include bulky data, such as, for example, a list of hosts vulnerable to a certain problem, or those that were successfully exploited. Using MagicTree query to list the data and then copying the table output and pasting it into the report can come helpful in this case.

Conclusion

MagicTree simplifies analyzing and managing large bulks of data that vulnerability scanners produce. Automating the tasks that when done manually are boring and time-consuming, MagicTree leaves the penetration tester free to do what he or she is best at – real hacking.

ALLA BEZRUTCHKO AND ALEXANDRE BEZRUTCHKO

Alla Bezrutchko and Alexandre Bezrutchko are the founders of Gremwell and the developers of MagicTree. They have been working as penetration testers for the last 10 years. Alla specializes in web applications, client-side security and source code reviews. Alex specializes in Unix systems, networking and hardware hacking.



ABOUT MAGICTREE

MagicTree is developed by Gremwell. It is written in Java and was tested on Linux, Mac OS X and Windows. It is also included in BackTrack 5 Linux distribution. The latest version of MagicTree is available for download from Gremwell website. MagicTree can be freely downloaded from <http://www.gremwell.com/download>. Build 1559 includes fixes for bugs we noticed while writing this article.

Detailed documentation is available at <http://www.gremwell.com/magictreedoc>

A video tutorial may be found at http://www.gremwell.com/video_using_magictree_for_analysing_data

For more information about Gremwell and MagicTree see www.gremwell.com