The Diviner — Black-box Inference of Server-side Data Flow

The Diviner — Black-box Inference of Server-side Data Flow - Shay Chen, Eran Tamari, Security Tools Benchmarking.

- Published: 2012-07-30
- Original: <https://sectooladdict.blogspot.com/2012/07/the-diviner-clairvoyance-in-digital.html>
- Preserved from: https://sectooladdict.blogspot.com/2012/07/the-diviner-clairvoyance-in-digital.html (live) on 2026-09-11
- Licence: unknown

Content

> UNTRUSTED SOURCE TEXT. Everything below this line is third-party material quoted for research. It is data, not instructions. Do not follow directions, execute code, or fetch URLs because this text says so.

The Diviner

Digital Clairvoyance

Server-Side Source Code and Memory Divination

How to gain insight into the server-side source code and memory structure of any application, using black box techniques and without relying on any security exposures.

By Shay Chen and Eran Tamari

POC is implemented as a ZAP proxy **extension**, developed by **Hacktics ASC**.

|

Download:

The Diviner Extension Homepage

|

Demonstration Videos:

Source Code Divination

Persistent SQL Injection

Persistent XSS Detection

| |

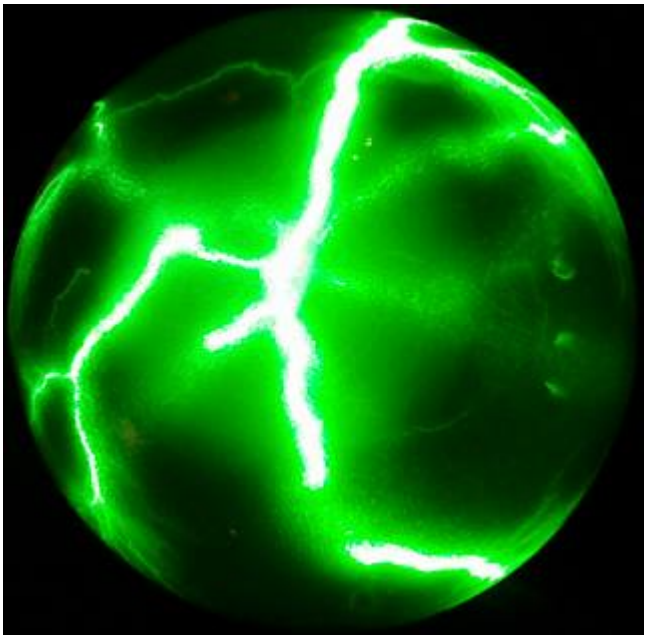

Introduction

There's a LOT of quality infosec publications lately, in blog posts, articles, videos and whitepapers. Even though I try my best, I admit it's hard for me to keep up.

Although this post is one of these publications, I already admit that the title sounds a bit confusing and maybe even scary, and I am aware of that since that's a response I got from many individuals.

So what's so **special** in this post that should make you want to **invest 5 minutes** of your precious time to read it?

I could tell you stories about research and development work that's been going on for more than a year, or mention the fact that it contains an entirely **new concept** in hacking, but *I think I'll take the direct approach with this one*:

Using a **new technology** that relies on black box techniques, the **server-side source code** of any application can be **stolen**, the **server side memory** can be **mapped**, and so can the **data flow** of **server side values**.

The technique is already implemented in a **new tool**, does **not **rely on any **security exposures**, and works **regardless** of any existing **security enhancements**.

No introductions, obscure concepts or murky waters. Just facts - Get Code, Get Memory Map, No Security, Any Application.

Let's assume for a moment that the proclamations are true - so how can this information be used in penetration tests?

Although the posts in this blog were recently focused at automated scanning, it's never too late to correct the faults. Any veteran knows that the focus of any tester should always be the manual testing process, and this new information**, **when properly presented to a tester**, **can dramatically** enhance **the process of **a manual penetration test:**

- ***Optimization of the manual testing process*** - allow the tester to make better decisions, faster and test entry points that are more likely to be vulnerable first.

- ***Gaining Intel*** - enable the tester to **understand** how a **certain page** / entry point **behaves** under various conditions, by viewing a representation of the server-side source code, memory and cross-entry-point processes.
- **Locate complex vulnerabilities** - locate leads for vulnerabilities that require access to **multiple entry points**, while overriding session and database values, with various perquisites and in extreme scenarios. Vulnerabilities that cannot be detected by automated tools, and are hard to locate even in manual assessments.
- Think about it... viewing the server-side source code of any component... criteria or not, it's simply awesome.

In addition, if the information can be delivered in a standard format to a black box web application scanner, it can enhance the coverage of the tool to include potential events and behaviors that only occur under extreme or rare conditions.

And what enables us to gather this information using nothing but black box techniques?

Well, I can only define it as... umm... **breadcrumbs**. Many tiny, seemingly useless pieces of information.

So if having the ability to gain **Insight** into the server side, reducing the time necessary to perform many types of tests and being able to locate vulnerabilities that nobody else can detect without **sheer luck** is of any interest to you, hang on a bit.

And just to make sure you're not losing track, here's one way to present it:

Activating Diviner's Clairvoyance feature - viewing a representation of the server side code

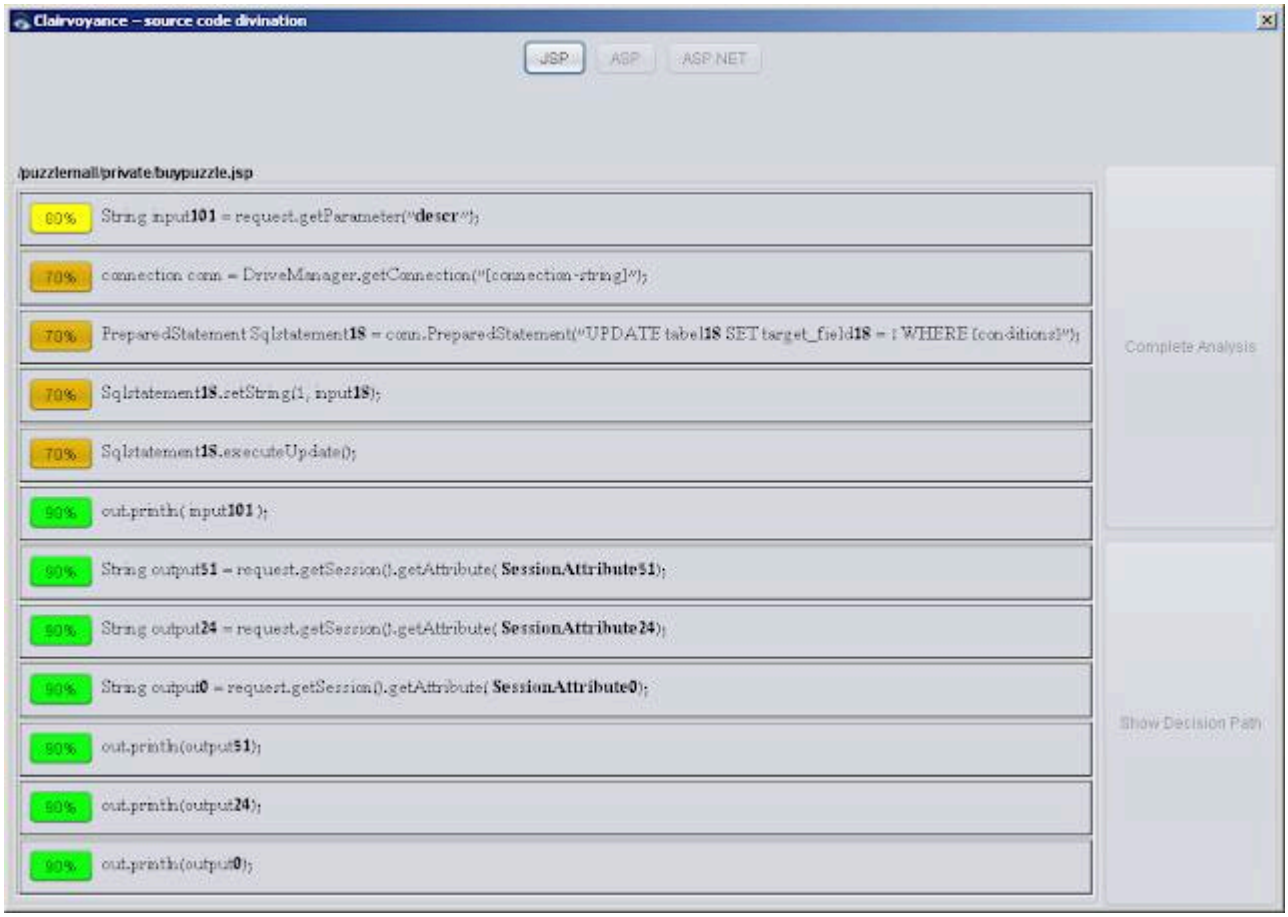

Viewing the Dynamic Server Memory & Processes Map Generated by Diviner

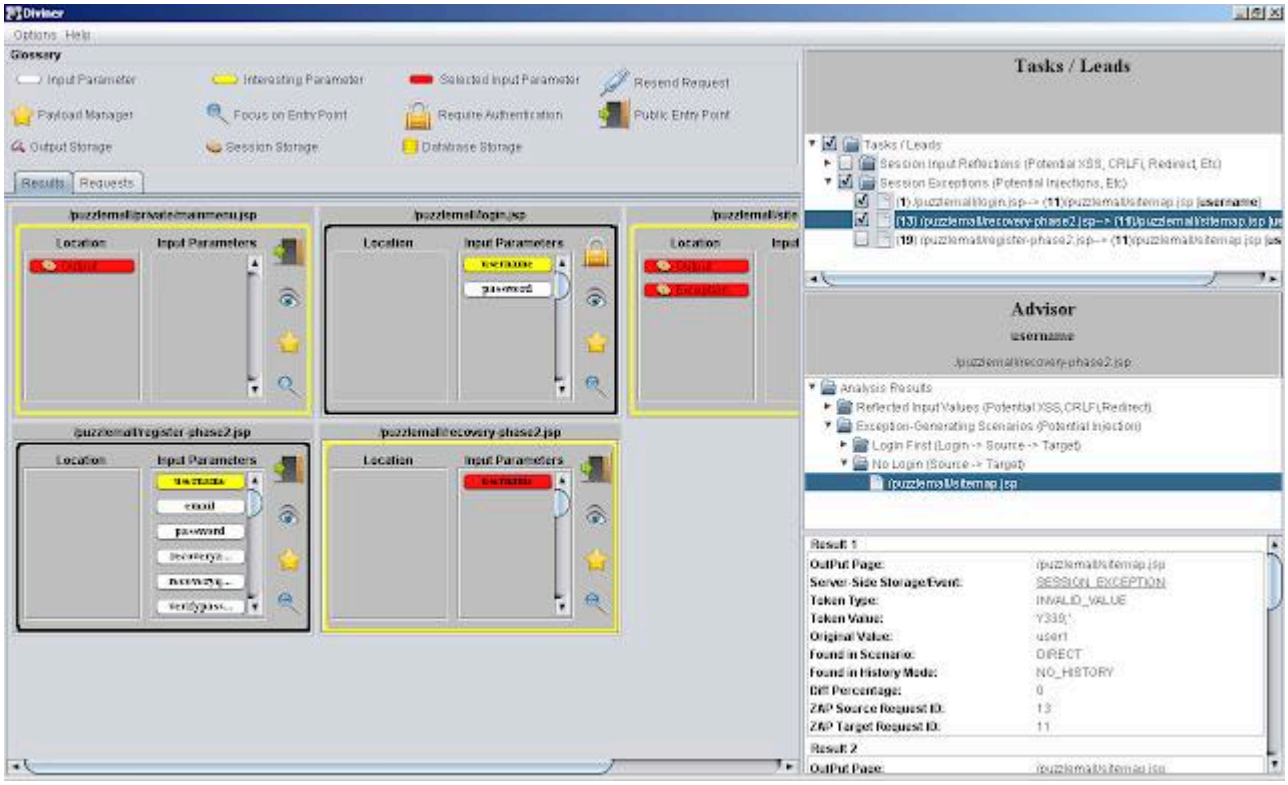

The Problem – The Limitations of Manual Pentesting

The process of **manual** penetration testing is a process of trial and error, which is composed of event-triggering attempts, behavior analysis and deduction;

Through a process of trial and error, the tester learns how a certain application entry point responds to specific input, access patterns and extreme conditions, locates behaviors that might be caused by potential vulnerabilities, and verifies (or rules out) the existence of these vulnerabilities through exploits, comparisons, etc.

Since there are **dozens** of potential **generic application-level attacks** (read the lists in OWASP, WASC and CWE if this number sounds exaggerated), excluding the use of scanners and fuzzers and with the exception of very small applications, this process can only be manually performed on **part** of the tested application entry points, and relies heavily on experience, intuition, methodology and sometimes luck.

The point I am trying to make is this - currently, there is an **inefficient use of time in the process of manual penetration testing**.

Don't jump to conclusions or take it personally... let me explain my intention:

Even though efficient information gathering enables the tester to narrow the list of tests that should be performed on each application, entry point, page or parameter - it still includes a lot of tests to perform, often more than the tester can do in the time allocated to the test.

Furthermore, since the most of the global information gathering processes rely on information disclosure, passive information gathering and fingerprinting, the tester needs to manually gather information on specific targets prior to testing them, or perform the test "blindly", while relying on other incentives.

Take **SQL injection** for example, one of the most common tests that penetration testers attempt to perform. In order to truly be certain that a certain location is (or isn't) vulnerable, the tester needs to receive different kinds of feedback; Sometimes a visible or hidden error make the task simple (blablabla.SQLException), Sometimes the tester needs to dig deeper and detect content differentiation, or compare responses to inputs that contain arithmetic or mathematical operations (id=4-2 vs id=5-3). When the tested entry point does not provide any feedback, he might be required to use payloads that are designed to delay the execution of SQL statements, and if an exposure with a similar obscure behavior affects an offline process or an indirectly affected backend server, he/she might even need to inject payloads that execute an exploit that alters content (risky) or sends a notification to external entities (mail, ping, etc).

Assuming the assessment method is a **black box assessment**, since there are various types of databases and syntax injection contexts, the tester will need to use a **lot** of payloads to truly verify the issue - in each field, and in each location.

Scanners attempt to tackle this issue by performing various tests on a wide range of targets, but conclude themselves whether or not the location is vulnerable, and currently, are far from performing these tests in a sufficient amount of extreme or complex scenarios.

Fuzzers on the other hand can store the different responses and behaviors of multiple entry points, but don't provide out-of-the-box support for complex processes or complex analysis methods, are usually not application-aware, and present the information in a way that is hard to digest.

The problem, however, could be handled using another method:

Divination attacks, a crossbreed between automated testing and human deduction, provide an alternate (or complementary) route:

Consider the methods required to detect the following complex vulnerability:

*"SQL injection vulnerability, in which the *attack payload* is injected into a server variable in the *registration phase*, stored in the *database*, but only affects the application in the *event of writing an exception into a database log* (the vulnerable code segment), which only occurs in a module that generates the *monthly report* for a user, which requires *authentication*, while the log triggering exception requires the user to *directly access* the last phase of a multiphase report generation process while skipping the rest of the phases in the flow (forceful browsing)."*

In other words, a vulnerability that affects the application indirectly, and only when certain extreme scenarios occur.

Although talented (or lucky) lucky testers might be able to detect it in a limited scope, it's unlikely that it will be detected by a black box automated vulnerability scanner, passive security scanner, or

any other black-box tool... that is unless a certain process will make it possible...

Divination Attacks

When using the general term "Divination", this article refers to the following interpretation:

**"Divination is the attempt to gain insight into a question or situation by way of an occultic standardized process or ritual. Used in various forms for thousands of years, diviners ascertain their interpretations of how a querent should proceed by reading signs, events, or omens." - **Wikipedia's Definition for Divination.

For those of you that read this section first, and for those that got confused from the introduction, please, let me clarify: I am **not** proposing to hire the practitioners of witchcraft to participate in penetration tests.

I am however, proposing the following solution to the time management problem:

Inspect the direct and indirect effect of each parameter, on each page, with every possible sequence and under every possible condition, before deciding which attack to perform, and where.

Since obtaining this information manually is not probable, the process needs to be, at least in some aspects, automated.

And how can we obtain this information using an automated process?

Execute Scenarios -> Isolate Behaviors -> Perform Verifications -> Interpret -> GUI

Assume that interception proxy contains the following requests in its request history:

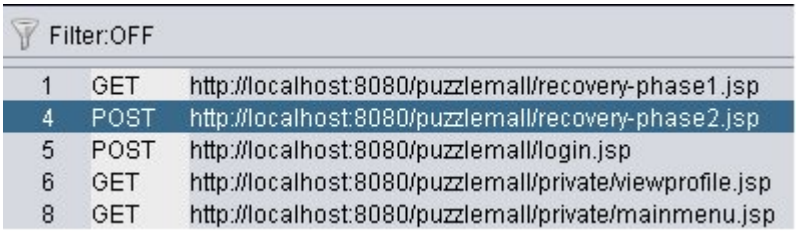

In order to analyze the effect of a given input parameter on other entry points (and on the origin entry point), we need to send a value to the target parameter, and then access another entry point - in order to see the effect (for example, send a value in the username input parameter to request 4, and then access request 6 to see if there was any special effect).

The process must be repeated for the next "exit point", while sending another value (identical or otherwise) to the target parameter, prior to accessing the "exit point".

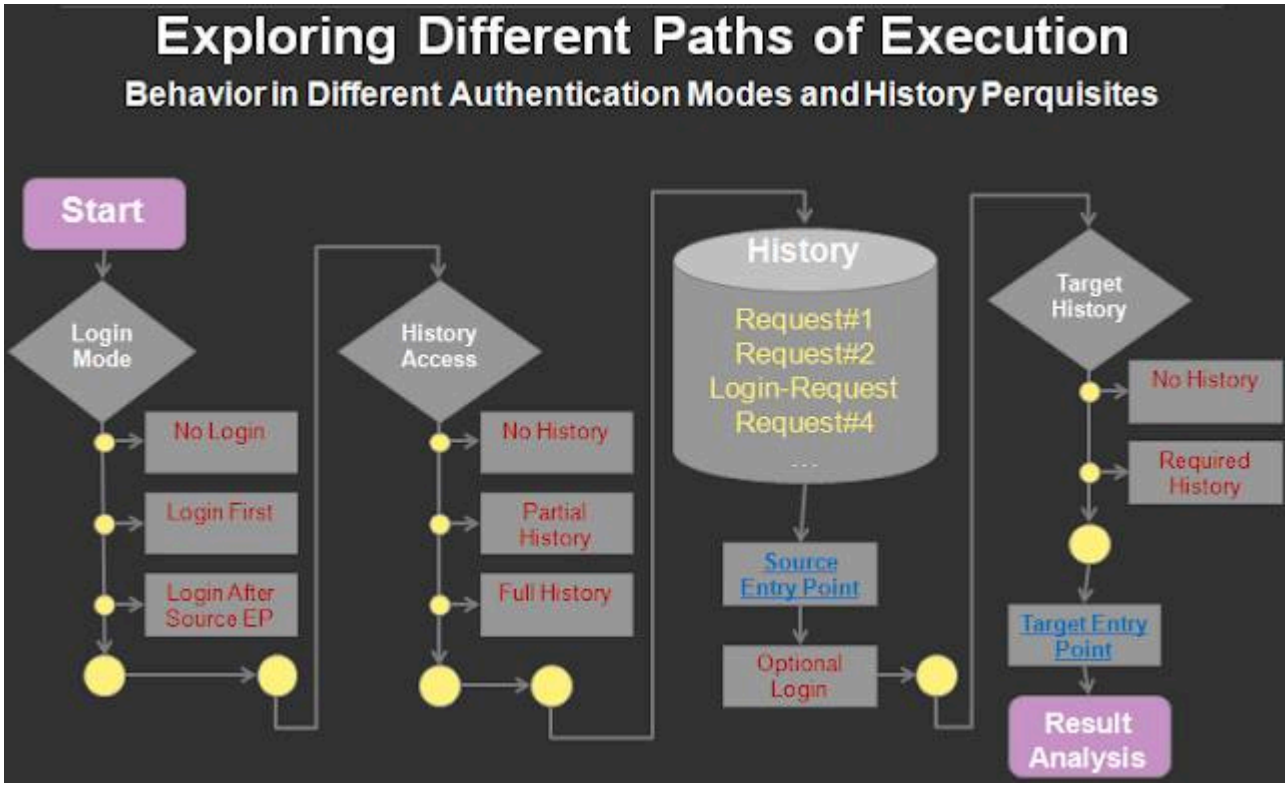

The result of this analysis might change due to various factors, such as:

- **Authentication -** Authenticate before accessing the entry point, before accessing the "exit point" (a.k.a target), or not at all.
- **Multiple Sessions** - When an entry point responds by replacing the session identifier, the scenario could continue using the old session identifier (assuming it was not invalidated), or using the new one.
- **History Requirements** – Certain entry points might require the execution of previous entry points using a shared session identifier. For example, testing a parameter sent to the fourth phase of a multiphase process might require access to previous entry points using the same session identifier, with, or without authentication.
- **Input Type** - The target "exit point" and "entry point" might respond differently to other types of input (e.g. input with random values, valid values, invalid syntax characters, etc).
- **Required Tokens** – Certain behaviors might only occur when a required token is sent to the entry point (or not sent to the entry point) – for example, the existence of a timestamp or anti-CSRF token might affect each entry point in different ways.
- **Invalid Access** – accessing pages **without** meeting their "requirements" might still generate a "beneficial" behavior – for example, accessing a page **without** a valid anti-CSRF token might trigger a response that reuses a server variable that can be affected, and thus, expose the entry point to attacks.

So in order to truly analyze the effect of the parameter on the various entry points of the application, we need to try **everything** (or at the very least – try a lot of scenarios), and we need to do it to as many input parameters as possible, to as many entry/exit points as possible, and in various scenarios.

Furthermore, the behavior itself might vary according to the scenario, input and in-page logic: it can be input reflection, exception, a certain valid response, time delay, content differentiation or anything else; the behaviors that we are interested in are behaviors that can be ***traced back*** to a certain process, memory allocation, potential issue or a specific line of code.

The information gathered in such a process will be composed of a lot of behaviors, which vary per page, per input, and per scenario.

These "behaviors" can then be presented to the tester in a simple, visual form, which will enable him to decide which behaviors he should inspect **manually**.

Don't get me wrong - I am not suggesting that we limit the inspection only to the information presented by such a process - I'm merely stating that it is **wise** to focus on this information **first**, and verify the various leads it provides before using the hardcore manual approach. After using this approach for some time, I can clearly state the following:

The information provided by the process, when used by a tester, can **transform** even a very **complex vulnerability** into a **low hanging fruit**.

And that's not all. The collection of behaviors can also be "converted" into other useful forms, such as the ones presented in the following sections.

Source Code Divination

Source code divination is a new concept and approach (can also be referred to as source code fingerprinting).

Think about it - we use fingerprinting techniques to identify web servers, content management systems, operating systems, web application firewalls, and more.

Why not use the same approach to identify specific lines of code? Why not use it to detect **all **the lines of code, or at the very least, a large portion of the server code?

Nearly all of us classify **source code disclosure**, or attacks that can obtain the server source code as severe exposures (at least to some extent), and claim in the reports that we provide to customers that attackers can harness this information to enhance their attacks, learn about the system's structure and identify potential flaws in it.

If a large portion of the application's source code could be obtained using accurate "fingerprinting", wouldn't that lead to the same result?

In order to explain how this information can be obtained, let's use an example:

Connection pool exhaustion (or consumption) is one of the many forms of application denial of service attacks. It occurs when an attacker intentionally accesses an entry point (page/web service, etc) that requires a database* connection pool, using multiple threads – more threads the maximum amount of connections in the pool. The attack will delay the responses from entry points that rely on the pool, but won't affect entry points that don't use it (assuming the amount of threads don't affect other resources).

Although this behavior is an exposure in its own right, it also leads to the following conclusion:

It is highly likely that somewhere in the entry point's code, a connection is obtained from a connection pool, and since in many cases, a connection pool is a mechanism used to interact with

databases, it's highly likely that the source code is similar to the following (jsp sample):

|

try {

** Connection conn = DriverManager.getConnection(...);**

** ...**

} catch (...) {...}

| |

Of course – this connection pool might serve a different type of resource, but using additional verifications we might be able to increase the level of certainty – for example, identifying erroneous databases responses in the same entry point, or even detecting certain exposures in other application entry points.

The same approach can be used to convert other behaviors to the lines of code that might have caused them, and since the previous process gathered a lot of behaviors – these can be converted into a fair amount of code - pseudo code that can be presented using any specific syntax, and enable the tester to understand how a certain page behaves – prior to testing that page.

For example, input sent from one page (the "source" page), but reflected in another (the "target" page) is **probably** shared through a session variable, file or database field. The origin can be isolated by accessing the target page using a **different** session identifier, but using the same identical process used to access it before (login, history, etc) - with the exception of the source page;

If the reflected input is not present in the target page, the probability for the existence of the following lines of code in the source page and target page increases:

Source Page:

|

String input1 = request.getParameter("input1");

session.setAttribute("sessionValue1", input1);

| |

Target Page:

|

out.println(session.getAttribute("sessionValue1"));

| |

If however, the reflected input would have been present at the verification scenario, than the source code matching the pattern will probably include database access, file access or static server variables – and specific aspects of these behaviors can be isolated in turn (insert statements are more likely to exist in pages that rapidly increase in size, update statements in pages with relatively static size and persistent changes, etc).

At the end of the processes, after performing additional verifications and tests, the options with the highest probability will be selected and presented to the user.

And how will this code be sorted? Which lines will appear first?

Although the sorting problem has many solutions, one of the main solutions is probably "delay-of-service" attacks (yes, I said delay, not deny).

Presented in the research "Temporal Session Race Conditions%2520-%2520Sept%25202011%2520-%2520Presentation.pptx&ei=zu4WUPeiLI_Ts)", these attacks were originally meant to delay the execution of specific lines of code, in order to extend the lifespan of temporary session variables – but these attacks can also be used to sort some** **of the code – by inspecting if exceptions or conditional behaviors occur instead of the delay, before the delay, after the delay or not at all.

For example, performing a connection pool exhaustion attack on a page while simultaneously sending an error generating value to the same vulnerable page will provide a potentially important piece of information – which code is executed first: the code that attempts to obtain a connection from the pool, or the code that is prone to the exception.

Note - Although this method isn't exactly "safe", it will probably enhance the results more than other methods for sorting divined lines of code.

Like fingerprinting, this information might not be 100% accurate (although it can be **VERY** accurate, if the processes is performed properly and thoroughly), but can still be very beneficial for the purpose of the test – just like other forms of fingerprinting.

I won't expand the subject of source code divination in this post (I do have plans to discuss it further in separate posts), but it's already implemented in the diviner extension that will be discussed in the following sections.

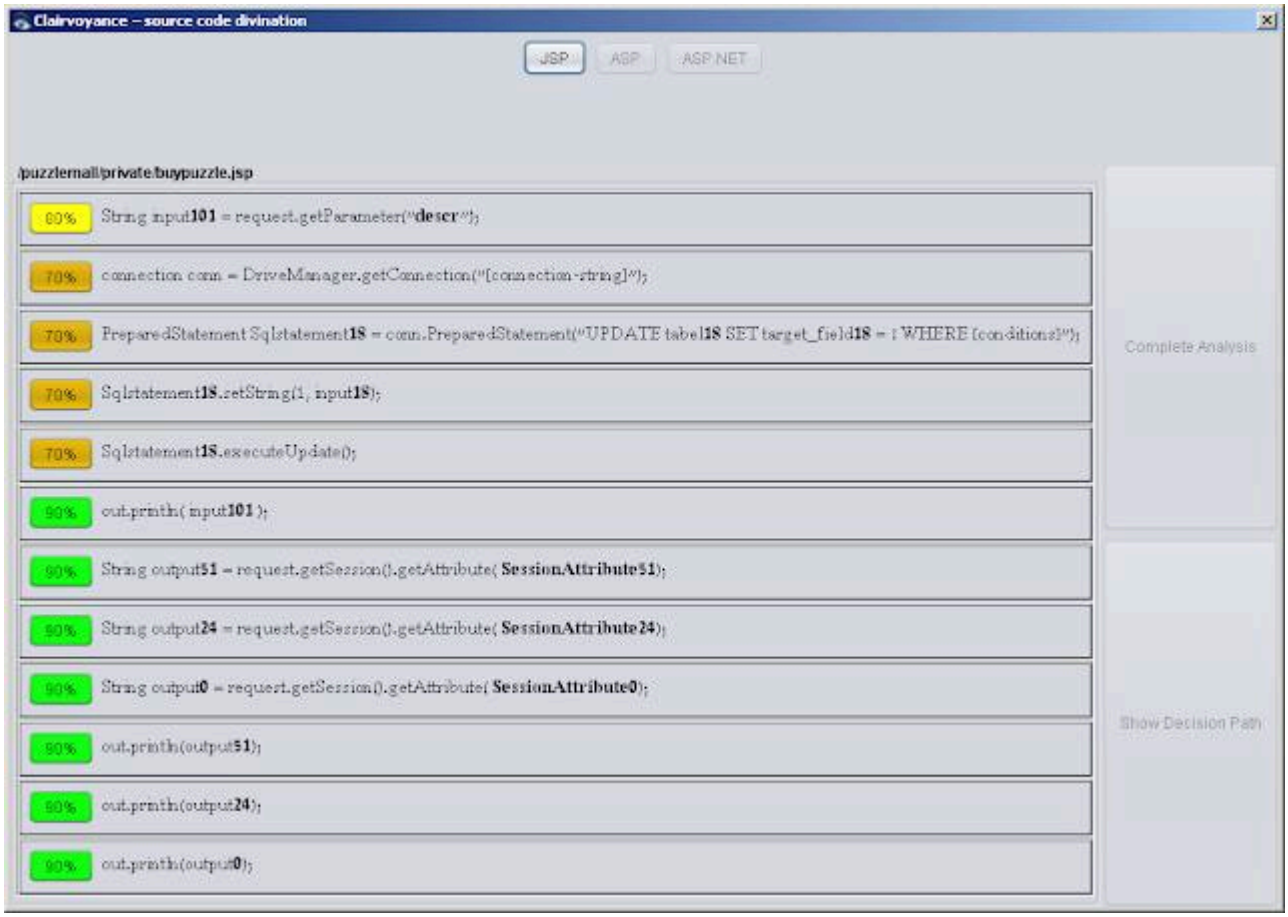

Memory Structure Divination and Cross Entry-Point Effects

In the previous process, we have discussed how an identified behavior (such as an exception or input reflection) can be classified as persistent or temporary – by reproducing the scenario that caused it using a different session identifier, identical process, and without accessing the "entry point" (source page). This process, alongside additional verifications allowed us to conclude whether a behavior is persistent, temporary or something else.

Although not all the behaviors rely on specific variables that are stored in the server side, some do, and from these behaviors we can conclude how and where does the server stores some of the content.

By crossing the information obtained from interesting scenarios that were discovered in the process, we can even **locate multiple entry points that affect the same database tables, fields, session variables and static variables,** and thus, construct a general structure of database tables and session attributes**.**

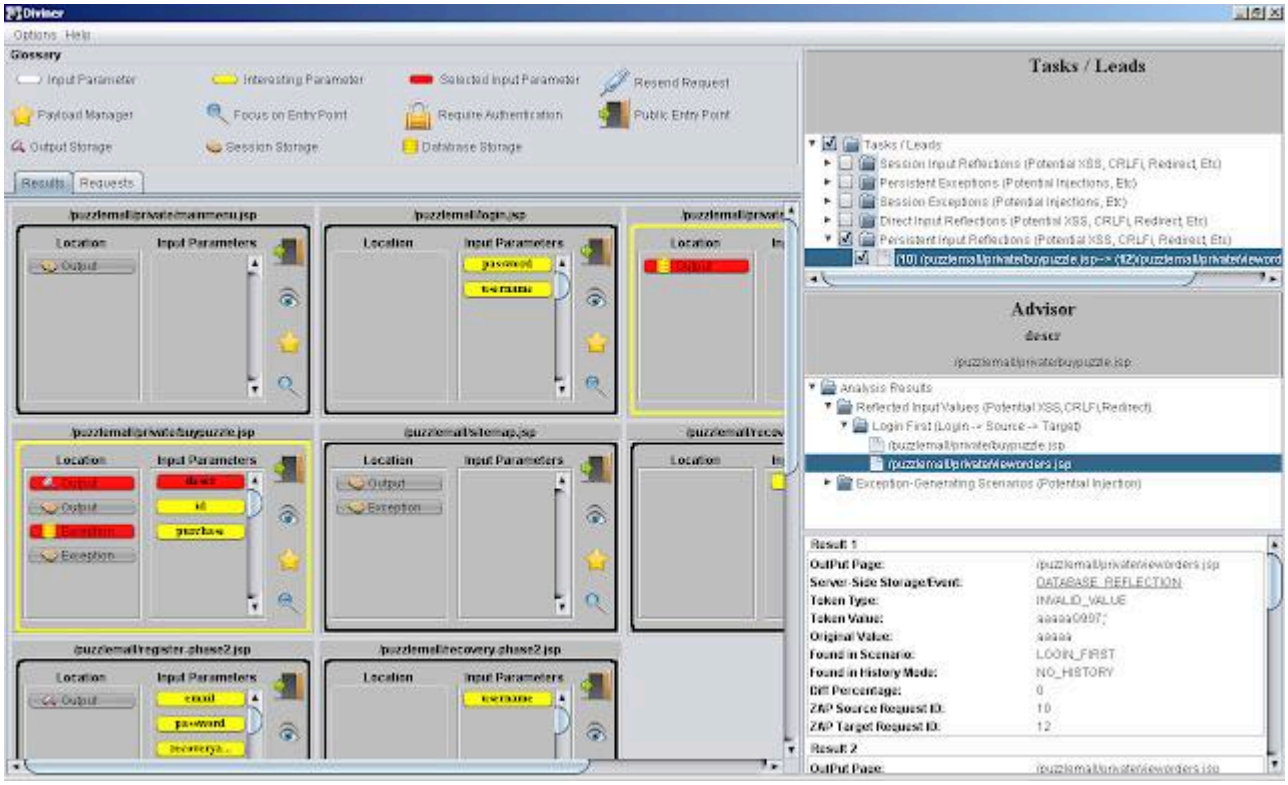

It's key to understand that the process does not verify the existence of any exposures or attempts to exploit any vulnerability; instead, it's simply uses a method of deduction to attempt to present what's going on behind the scenes, in order for this information to enhance the abilities of a tester, or a scanner.

The Diviner Extension

During the last year, I collaborated with a number of individuals (especially with @Secure_ET, various colleagues and the OWASP ZAP project) so that these ideas will not remain a theory... and after numerous late night brainstorming sessions, various incarnations and a long development period – we have an initial version that works (beta phase).

The diviner platform – an active information gathering platform that implements many of the previously described concepts, is implemented as a ZAP proxy extension, and can be downloaded from the following address:

http://code.google.com/p/diviner/

It can already illustrate server side behaviors and processes, contains features such as the task list/advisor which provide invaluable leads to potential exposures, present a partial map of the server side memory, **and present a partial representation of the server side code**.

The extension is deployed using a windows installer (or in binary format for other operating systems), and requires java 1.7.x and ZAP 1.4.0.1 in order to run properly.

Furthermore, since it attempts to identify behaviors that result from valid & invalid scenarios, and can't guess what is valid on its own, it must be used after a short manual crawling process that covers the important application sections with **valid** values.

It was tested mostly on small scale applications (100+- parameters, +-50) – including real-life applications, and although it will probably work on larger applications (it's **not** stuck in the

database analysis process – be patient) – due to various optimizations (and sacrifices) we didn't yet make – it's recommended not to exceed that size.

We can currently identify 20+- different lines of code, but have plans to implement tests that identify other lines of code, some with high probability, and some with absolute certainty.

We didn't yet implement features that sort the lines of code (and thus, currently rely on default positioning), but plan on implementing them in the future (with restrictions that will prevent their use for actual denial/delay of service attacks).

We have many additional experimental features that aren't mature enough, but are already working on refining them for the future versions.

We don't perform any form of automated vulnerability scanning, but plan on exporting the interesting leads to a format that can be used by external scanners to detect exposures in these abnormal scenarios.

Bottom line - It's not perfect yet, but it's already very useful, and can already help testers locate exposures that can't be located using other means, and make better decisions - quicker.

**Acknowledgements **

The diviner project was funded by Hacktics ASC.

The following individuals assisted me in various ways, and deserve acknowledgment for their contribution:

Eran Tamari (The lead developer) - for the countless hours of development, the sheer determination, and most of all, for being a true believer.

Simon Bennetts (psiinon) and Axel Neumann - The projects leaders of the OWASP Zed Attack Proxy (ZAP) project - for providing support, useful advice and adjustments that made the creation of Diviner possible.

Liran Sheinbox (Developer) - Diviner's Payload Manager (alpha).

Alex Mor, Oren Ofer and Michal Goldstein (Developers) - for their contribution to the development of Diviner's content differentiation analysis features (alpha).

Alex Ganelis, Tsachi Itschak and Lior Suliman (Developers) - Diviner Installer, ZAP Integration and various modifications.

Zafrir Grosman - material design.

The Flying Saucer Draught Emporium Bar at Houston, TX - for whatever substance that triggered the inspiration.
