OWASP AppSensor - CrossTalk

Article submission for the "Protecting Against Predatory Practices" themed September/October 2011 edition of CrossTalk, The Journal of Defense Software Engineering.

Draft v0.9 (1st April 2011)

Title

Building Real-Time Defenses into Software Applications

Abstract

Attack-aware software applications respond in real time to suspicious and malicious events with a very low false positive detection rate. The approach is especially suited to software applications with high information assurance requirements such as in the defense, critical national infrastructure and financial service sectors to protect against cyber espionage, fraud, tampering and theft. The OWASP AppSensor Project has developed a methodology, documentation, example code and pilot demonstration for this application-specific attack detection and response.

Copyright

Building Real-Time Defenses into Software Applications

Introduction

Information systems and data are being targeted relentlessly by skilled and motivated adversaries who are well resourced and have excellent tools. The attackers may be backed by organized crime, governments or commercial enterprises, and attempt to access sensitive and secret data available to software applications by finding and exploiting vulnerabilities in the applications themselves.

In this article we discuss why conventional application defenses are not a solution to these types of advanced attack, and explain an initiative from the Open Web Application Security Project (OWASP)[1] which utilises self-defence to provide real-time detection and response.

Conventional Defensive Measures

A fundamental starting point for secure software applications is to build security in at all stages of the development life-cycle. We need robust code designed, developed, configured and operated on hardened operating systems. This does not mean the software is impregnable; unknown vulnerabilities almost certainly exist, new vulnerabilities are introduced with software code changes and attackers examine and probe applications to find theses. Once they have identified one or more vulnerabilities, they need to discover ways to exploit them individually, or in combination. It is very hard to detect these custom attacks using conventional defensive measures. One reason is that conventional defenses either do not work at the application layer or they are designed with standard signatures to look for known common bad patterns. This may detect generic attacks, but will almost certainly fail to detect a custom attack looking to exploit a weakness within the application.

This genericity of conventional defensive measures means they can be applied to numerous applications, often with little to no tuning, with obvious benefits (cost, time to deploy, etc.). However, a significant challenge is a lack of context. Context is what allows applications to have intimate knowledge of both data and the users accessing that data.

Some techniques often thought to defend applications include using transport layer security (TLS/SSL), network firewalls, application gateways (e.g. web application firewalls, XML appliances, cross domain guards) and network/host intrusion detection and prevention systems (IDPS).

TLS is not a complete solution to the services required for a secure application. In addition to confidentiality and availability, minimum security standards dictate some kind of data integrity, along with authentication and authorization. TLS provides some degree of data confidentiality and integrity in transit and limited assurance in the identity of one or more of the parties depending upon how it has been configured. Most deployments use a single certificate model as e-commerce applications often perform their identity checks at payment; this however subverts a very important design criteria of the TLS standard which derives it security strength from a dual certificate exchange. However there is no inspection of the data itself — it simply passes on what was sent. Attackers' payloads are sent to the server in the same way as normal traffic. In the end, since TLS simply encrypts attackers' malicious data as it is sent to the server, it provides no protection to prevent an attacker from attacking the website.

Network firewalls are a vital component for network defense, but they allow or deny traffic using a particular port. By necessity web applications need to allow traffic through specified ports (often TCP ports 80 or 443). Attackers of course send their payloads using the same ports. And it is for this very reason, that it is commonly believed that the perimeter defence model is dead. Application gateways are generic solutions to typical problems. For example, even with careful configuration and self-learning, most web application firewalls are operated in detection-only mode, and they have no insight into business logic type of attacks. Like network firewalls, IDPS has no real concept of good and bad application usage.

Additionally, some of these conventional defenses try to replicate part of the application logic, but they can, at best, only do this partially. This external logic then needs to be verified, maintained, protected

and managed, adding another often significant overhead to operational costs.

In today's world, our applications are the gateways to our data and attackers are constantly probing to determine weaknesses in our systems. This needs a different approach than these conventional defenses.

Application Defensive Measures

The right way to detect advanced attacks is within the applications themselves, by adding application-specific security controls which detect malicious activity.

Consider the analogy of control instrumentation in an industrial production process. Measurement points for temperature, pressure, mass, velocity and volumetric, etc are added to quantify the state of key aspects in real-time. This information may be used in localised control (e.g. a float switch in a liquid storage vessel), but more often is integrated into an overall process control system, which aggregates signals from many sensors, and uses feed-back and feed-forward control mechanisms to react to changes and maintain a desired state.

Unfortunately, most current software applications are like industrial processes without this control instrumentation — there are no sensors and no intelligence about what is going on — they are blind to what is happening internally. Yet applications have access to powerful processing and storage capabilities in their supporting hardware and software systems.

Software applications have full knowledge about the business logic and the roles & permissions of users — each application can make informed decisions about mis-use, and to identify and stop attackers with a very low false positive rate. An additional benefit is that this context-aware analysis occurs in real-time. This is an application-specific approach, not a generic one, because the controls are related and integrated within the application, and can be selected based on an assessment of risks specific to the process and data.

Similar ideas already exist in some software, but these are usually implemented as isolated processes and some may be undertaken reactively to events, or performed largely in a manual way. The concept described in this article focuses and formalizes this approach. It is about implementing proactive measures to add instrumentation and controls directly into an application in advance so that all these events (and more) are centrally analysed and responded to.

Attack-Aware Detection

Once sensors have been deployed within the application code (and possibly external sensors), detection and response can take place. The sensors themselves send any events they detect to the analysis engine. The engine is then responsible for evaluation of individual events and determination of an attack. Once an attack has been recognized, the analysis engine determines an appropriate response and executes that response.

Software applications then become attack-aware and thus it is possible to undertake proactive defense against currently unknown threats. Applications are the appropriate place to detect many types of attacks given the wealth of information available to the application.

Normal and Malicious Behavior

The approach relies on being able to differentiate between normal and malicious behavior. That is one of the biggest hurdles when security controls (attack detection mechanisms) are added outside the application; it is difficult to distinguish user intent and this leads to a higher rate of false positives.

The ability to differentiate between normal behavior, suspicious behavior and attacks (and to do it with a high degree of granularity) can also be an improvement in the usability of human interfaces of the software applications. The application knows which inputs allow keyboard input, which have client-side validation, which should not be altered by a user and which are from trusted systems. A conventional defense such as a gateway device applying a whitelist derived by auto-learning could block a password which has a trailing line feed character appended. The application might consider this invalid, but not

consider it suspicious since the extra character could be the result of a straightforward copy & paste action. Thus the application does not block normal usage unnecessarily.

Evasion and Unknown Attacks

A common attack technique is to use obfuscation to bypass black list inspection performed generic application defenses outside of the application. Another benefit of integrating attack detection within the application is that all data has been fully canonicalised. At this point it is much easier for the application to inspect the user submitted data and determine if the data is malicious.

In addition, detection within the application allows insight into unknown attacks. Many detection points are designed to capture activity that can only occur outside the normal flow of the application. In other words, a user could never accidentally trigger the detection point (e.g. a web page receives a POST message but is only designed for GET). In these scenarios, the detection points can provide alerts to odd behavior that could represent attackers probing for a new type of weakness within the application.

AppSensor Project

The concept described above was initially developed under the OWASP Season of Code 2008[2] by Michael Coates. The output was a document defining the AppSensor[3] conceptual framework, available as a printed book[4] and free PDF download[5]. Over the last two years a reference implementation in Java[6], further guidance, demonstration pilot[7], presentations and videos[8] have been created to support the initiative.

Like in an industrial process control systems, a whole range of sensors can be considered for incorporation by the instrumentation engineer. Most of the equivalent application sensors are located within the application code itself, but AppSensor also has the notion of inputs from other "external" devices. As an example, the software could alter its security posture based on an external threat level rating such as the United States Armed Forces Defense Readiness Condition (DEFCON).

AppSensor has defined over 50 detection points[9] (Table 1) reflecting the aspect being monitored. The detection points range from signature type detection points which typically look for particular events within a user's single request/response cycle, to behavioral type detection points which assess events over longer time periods (e.g. a single user session, a finite time period, the life of the application). The list is not definitive; but is a suggested starting point since custom applications will need to consider carefully adding detection points specific to the application at various points deemed important to the custom process logic within the application.

Table 1: AppSensor guidance lists over 50 example detection point definitions grouped into these 12 categories; additional custom detection points are also often required.

Detection Point Type	Code	Name
Signature	RE	Request Exceptions
	AE	Authentication Exceptions
	SE	Session Exceptions
	ACE	Access Control Exceptions
	IE	Input Exceptions
	EE	Encoding Exceptions
	CIE	Command Injection Exceptions
	FIO	File IO Exceptions
	HT	Honey Trap
Behavioral	UTE	User Trend Exceptions
	STE	System Trend Exceptions
	RP	Reputation

Although AppSensor works best within the authenticated portion of an application, where the user's identity is known and can be blocked if malicious, it is also possible to apply the principles to other areas.

Documentation has guidance on detection point considerations, determining the malicious intent, monitoring system trend events and implementation guidance. However, the planning stages are probably the most time-consuming aspect of implementing AppSensor. If threat modelling is already being used in your software development life-cycle, this may be a natural location to determine many of your application specific detection points but recently additional guidance[10] has been created to assist with this stage.

Rich Responses

Once the detection points have been determined, thresholds are set for each detection point, or group of detection points, which then act like tripwires. In OWASP AppSensor, an analysis engine monitors the detection events and determines when, and what, response is appropriate according to these specified event thresholds.

Conventional defenses offer binary response options such as allow/deny or log/block (Table 2). They are often limited to a single dimension — the particular network connection contravening a rule.

Table 2: Conventional defensive measures may only offer a binary choice such as logging or blocking.

Response Type	Examples
Allow	Log events only
Deny	Block IP address

The AppSensor concept describes a wide spectrum of possible responses[11], and when combined with the very low false positive detection rate, defenses do not need to be configured in detection-only mode for fear of preventing valid traffic. Inside the application provides a two-dimensional view, where responses may be directed not only against a single connection or a single user, but now also against a whole group (e.g. based on location, role, or behavior), or against all users. Table 3 lists the most common responses, from those hidden from the user, to those which the user may be aware of and those which actually interfere with an action.

Table 3: AppSensor provides the ability to have a much richer range of responses; the guidance lists thirteen types of response. Specific examples for each type show how the types can be customized for each application's context.

Response Type	Examples
Logging Change	Full stack trace of error messages logged
	Record DNS data on user's IP address
Administrator Notification	Visual indicator displayed on an application monitoring dashboard
	Audible alarm in the control room
Other Notification	Signal sent to upstream network firewall, application firewall (e.g. XML, web) or load balancer
	Alert sent to fraud protection department
User Status Change	Increase data validation strictness for all form submissions by this citizen
	Reduce the number of failed authentication attempts allowed before the user's account is locked
Proxy	Requests from the user invisibly (from the user's perspective) passed through to a hardened system
	Request are proxied to a special honeypot system which closely mimics or has identical user functionality
User Notification	On-screen message about data validation issues
	Message sent by email to the registered email address to inform them their password has been changed
Timing Change	File upload process duration extended artificially
	Add fixed time delay into every response
Process Terminated	Discard data, display message and force user to begin business process from start
	Redirection of an unauthenticated user to the log-in page
Function Amended	Limit on feature usage rate imposed
	Additional validation requirements
Function Disabled	Web service inactivated
	Content syndication stopped

Account Logout	Session terminated and user redirected to logged-out message page
	Session terminated only (no redirect)
Account Lockout	User account locked permanently until an Administrator resets it
	One user's IP address range blocked
Application Disabled	Website shut down and replaced with temporary static page
	Application taken offline
Collect Data from User	Deploy a Java applet to collect remote IP address
	Deploy JavaScript to collect information about the user's network

When behavioral aspects over time (a request/response, a session, a time period) are included, the application becomes capable of rich three-dimensional responses. The behavioral aspects can span beyond application restarts, upgrades and even the application life cycles from deployment to disposal, since data from other systems which pre-date the application may be incorporated.

Cross Integration

As mentioned above, AppSensor can potentially make better use of information from other systems and devices to contribute to its pool of information for attack detection, increasing the value of those other systems. Other systems may be able to contribute to AppSensor's decision making processes. For example behavioral information can be fed into stochastic systems for statistical analysis, which may be useful in identifying user error from malicious behavior. Something that occurs very infrequently but causes an exception to be thrown may actually be down to human error, such as a typo, or other anomaly, whereas hundreds of the same event happening at once may indicate a malicious behavior. Hundreds of people are very unlikely to be making the very same error at the very same time. Additionally, this further reduces false positives since stochastic methods 'tune' the system for outliers.

As well as consuming data, AppSensor can pass back intelligence to other systems. This might be used as part of the real-time response:

- locking a user's single-sign-on account

- removing some permissions for a user from a physical access control system

- setting a warning flag about a customer in a CRM

- blocking a session ID or IP address at a network firewall.

Equally, data can be federated to other systems such as SIEM, IDS and fraud monitoring systems.

Information Insight and Value

Applications contain considerable knowledge about users and their activities. Without AppSensor so much intelligence is being wasted. System owners are left with tools and devices which have to guess about what is going on. AppSensor provides insight into whether, and how, attacks are occurring.

A pilot implementation has also demonstrated the value this approach has to containing the damage caused by application worms. System trend detection points can be used to detect large changes in a function's usage, alert administrators and ultimately disable the function. While this does not remove the infection, it stops the worm from spreading, limits the extent of data requiring clean-up, and may allow the remainder of the application to continue functioning, which may otherwise have had to be stopped.

And of course if this is a public-facing commercial application, some of these data would be valuable to those interested in understanding and improving operational objectives relating to metrics such as those for availability, usability, reducing drop-out rates, conversion ratios, etc. Another key benefit, is that this give security useful metrics, currently in absentia.

Applicability

Software applications must be built securely in the first place. If issues like authentication, session management, authorization, etc are broken already, it will not be possible to implement the approach.

Self-defense goes beyond the implementation of security services, but there does appear to be growing interest in this area. For example, the U.S. Defense Department recently announced[12] that it is planning to spend $500 million to research new cyber security technologies including "active defenses" - technologies that detect attacks and probes as they occur.

OWASP AppSensor is a comprehensive proactive approach, rather than reactive, which can be applied to applications in many situations. It reduces the risk of unknown vulnerabilities being exploited by identifying and containing users conducting malicious activity that is often the precursor to an attack. It greatly increases the visibility of suspicious events and actual attacks. This can provide additional information assurance benefits:

- reduced security risks to data and information systems
- improved compliance
- reduction in the consequences of data breaches.

In turn, these can provide improved service levels and resilience which have wide applicability. Some types of organizations may also achieve competitive advantage.

Conclusions

All types of organizations have very powerful systems with access to mission critical data, but generally have no idea if, when or how applications are under attack. This is the crux of the problem; critical data on critical systems and little to no visibility of the health of the overall environment. This must change. The concepts embodied by OWASP AppSensor present a solution to this issue.

Proactively detecting attackers via attack-aware applications can significantly increase the assurance of our systems. The AppSensor concept says that the application must also defend itself, as well as provide all of its own security services. By creating applications that are attack-aware and able to respond appropriately, the assurance of our systems is increased. In addition, there is better operational visibility of the applications that hold our most critical data.

Careful planning is needed to ensure the very low false positive detection rate is maintained, and that authenticated users are not prevented from undertaking valid tasks. Although AppSensor provides a low-risk method for incorporating detection points and responses using the ESAPI library, adopters who develop their own implementations need to consider the usual risks of custom-developed software. AppSensor is a practical approach which has been implemented in live software applications.   The AppSensor team are considering ways to improve its analysis engine including the use of stochastic methods, integration with other systems such as SIEM and building the concepts into software frameworks and libraries so the overhead is reduced. The team would also like to investigate how AppSensor could be applied in an Agile / Test Driven Development environment.

However, when combined with stochastic methods, we believe OWASP AppSensor can be the single most important pattern against unknown attacks to date. This defense against the unknown attacks is a giant benefit that can not be achieved using anything else, anywhere, at any price.

Acknowledgements

About the Authors

Colin Watson

Colin Watson's work involves the management of application risk, building security and privacy into systems development and keeping abreast of relevant international legislation and standards. He holds a BSc in Chemical Engineering from Heriot-Watt University in Edinburgh, and an MSc in Computation from the University of Oxford. He contributes to a number of OWASP projects and is a member of the OWASP Global Industry Committee. Colin is a consultant and co-founder of Watson Hall Ltd.

Email colin.watson@owasp.org

Michael Coates

Michael Coates has extensive experience in application security, security code review and penetration assessments. He holds an MSc in Computer Security from DePaul University and a BSc in Computer Science from the University of Illinois. Michael is the creator and leader of the AppSensor project, a contributor to the OWASP Top 10 and a frequent speaker at security conferences. As web security lead at Mozilla, Michael protects web applications used by millions of users each day.

Email michael.coates@owasp.org

John Melton

John Melton is a software architect and designer with particular knowledge of security for online systems. He is the AppSensor project's Development Lead, and has been responsible for most of the example code developed. He holds an MSc in Information Security and a BSc in Computer Science, both from the University of North Carolina at Charlotte. John is a Senior Information Security Engineer at Wells Fargo and lectures at UNC Charlotte.

Email john.melton@owasp.org

Dennis Groves

Dennis Groves is the co-founder of OWASP. He is a well known thought leader in application security who's work focuses on multidisciplinary approaches to information security risk management. He holds an MSc in Information Security from the University of Royal Holloway, University of London. He is currently a UK expert for the UK mirror of ISO subcommittee 27, WG4.

Email dennis.groves@owasp.org

1 The Open Web Application Security Project (OWASP). About the Open Web Application Security Project. <http://www.owasp.org/index.php/About_OWASP>

2 Ibid. Summer of Code 2008. <http://www.owasp.org/index.php/OWASP_Summer_of_Code_2008>

3 Ibid. AppSensor Project, <http://www.owasp.org/index.php/Category:OWASP_AppSensor_Project

4 Coates, Michael. AppSensor - Detect and Respond to Attacks from Within the Application, v1.1, OWASP. Paperback, 48 pages, Lulu. <http://www.lulu.com/product/paperback/owasp-appsensor/4520003>

5 Ibid. AppSensor - Detect and Respond to Attacks from Within the Application, v1.1, OWASP. <https://www.owasp.org/images/2/2f/OWASP_AppSensor_Beta_1.1.pdf>

6 OWASP. AppSensor Developer Guide. <http://www.owasp.org/index.php/AppSensor_Developer_Guide>

7 Ibid. AppSensor Live Tutorial. <http://www.defendtheapp.com/>

8 Ibid. AppSensor Media. <http://www.owasp.org/index.php/OWASP_AppSensor_Project#tab=Media>

9 Ibid. AppSensor Detection Points. <http://www.owasp.org/index.php/AppSensor_DetectionPoints>

10 Watson, Colin. AppSensor Implementation Planning Workbook, OWASP, December 2010. <http://www.owasp.org/index.php/File:Appsensor-planning.zip>

11 OWASP. AppSensor Response Actions. <http://www.owasp.org/index.php/AppSensor_ResponseActions>

12 Ratnam, Gopal and King, Rachael. "Pentagon Seeks $500 Million for Cyber Technologies", Bloomberg, 15 February 2011. <http://www.bloomberg.com/news/2011-02-15/pentagon-seeks-500-million-for-cyber-research-cloud-computing.html>