AdNovum

NOTITIA

Noteworthy news from and about AdNovum

Summer 2015, Issue No. 28

Make Security an Asset

Content

Dear Reader,

You may think poodles are cute. Be warned: This poodle belongs to the same family as FREAK, Shellshock, Heartbleed and BEAST. All of them originate from the virtual world ... and pose a very real threat. They are also known as vulnerabilities. Adequate protection for IT systems and applications is therefore highly advisable. Of course security always requires an investment as well. Taking the right steps improves the result – with genuine added value in the form of a competitive advantage. Read on to find out how to achieve this advantage and what else you should know about web security.

Stephan Schweizer, Nevis Product Manager, and Thomas Zweifel, Senior IT Consultant, take us on a journey through time in the introductory article. We visit the classic Internet age in the years around the turn of the millennium, when the fronts between good and bad on the Web were still clearly defined. Then we explore current technology, which opens up unimagined possibilities for us with HTML5. Yet it does the same for potential attackers.

In the interview, Gion Sialm of the Federal Office of Information Technology, Systems and Telecommunication (FOITT), Marc Condrau of Health Info Net (HIN) and Ivan Buetler of Compass Security discuss web security from the perspective of the federal government, an industrial enterprise and a security firm. They impart a realistic big picture from three different expert perspectives.

So how can an individual IT company such as AdNovum protect itself and its customers against virtual adversity? By integrating security into software development for example. In the background article Marcel Vinzens, Deputy Chief Technology Officer, and René Rehmann, Security Officer, report on what this means and which security processes have been implemented by AdNovum in concrete terms.

Our guest author Martin Kuppinger, founder and Head Analyst at KuppingerCole, dedicates himself to authentication as a central aspect of security. Since not only people but also apps and things access information today, secure access is more important than ever. Learn what type of authentication is recommended when and why.

Enjoy reading!

Chris Tanner

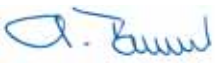

CEO AdNovum Informatik AG

Yesterday and today – the transformation of security solutions

A whole new generation of web applications is establishing itself with the success of JavaScript and HTML5. Applications are being networked across organizations at the same time. What does that mean for proven security concepts?

by Stephan Schweizer and Thomas Zweifel

In the classic Internet age at the end of the 1990s and early 2000s the fronts were clearly defined: The "bad" populated the Internet, the "good" did their work on the intranet. What is known as a network perimeter was and is being set up to separate these two worlds. This usually consists of an external and an internal firewall. Additional security components such as secure reverse proxies and web application firewalls are located in the zone in between them, known as the "demilitarized zone" (DMZ). They ensure that only users with strong authentication (through what is known as two-factor authentication) can access the available web applications from the "evil empire". This barrier of security components also prevents direct communication between the client and application. In addition to single sign-on functionality, such a setup also offers flexible options for intervention in order to protect applications against direct attacks such as denial-of-service (DoS), cross-site scripting and injection attacks (see box).

Around the turn of the millennium, the fronts between good and evil on the Internet were still clearly defined.

The application logic is shifting to the client

In this world there were no doubts: The application logic lies on the server side. Providing visual access to content for the user was the primary task of the browser. There was a broad consensus that this was the right approach after the experiences made in the client-server era of the 1990s. Users were very frugal as well, being satisfied with spartan user interfaces (GUIs) that appear archaic from today's perspective.

Denial-of-service (DoS) attacks *are aimed at making the use of a service by regular users impossible. The attacker attempts to overwhelm the service with a large number of requests, or to make it crash by purposefully exploiting known vulnerabilities.*

With distributed denial-of-service (DDoS) attacks, *a large number of Internet clients distributed around the world is used to overload the target system with requests. These clients are usually workstations that were previously infected with a Trojan virus. Subsequently, the system is remotely controlled by the attackers so the owners usually participate in the denial-of-service attack without their knowledge. The large number of sometimes varying clients makes it more difficult for the service provider to filter the DoS requests.*

Cross-site scripting *aims to infiltrate a web application with the attacker's JavaScript code. The attack usually takes place via insufficiently protected input fields of the application or URL query parameters. Depending on the attacker's intentions, the code that is smuggled in uses the browser's runtime environment to intervene in the application logic, load additional malicious code, steal digital identities or exploit existing security vulnerabilities on the client system, for example to install a Trojan virus.*

SQL injection *attacks are aimed at spying out, stealing or modifying data (such as credit card information) stored in application databases. They attempt to inject and execute the attacker's SQL statements in the application through unprotected input fields or query parameters. SQL Injection attacks can be prevented through corresponding WAF filters or on the database side also through "Prepared Statements" (with bound parameters).*

GUI requirements became more elaborate with the development of superior applications. Initial attempts aimed at making GUIs user-friendly were made with JavaScript. At that time however, it would not have occurred to anyone to implement business logic in JavaScript – the elements of an application realized in JavaScript were limited to logic related to visualization. There was good reason why developers were reluctant to implement functionality in the browser: In the former HTML4 world, the tags may have been standardized but the semantics and representation were interpreted very differently by various browser providers. This meant that Web content looked entirely different depending on the browser that was used. The JavaScript API (Application Programming Interface) provided by the browsers was very spartan as well. Those who wanted to tackle more sophisticated tasks with JavaScript while simultaneously gaining control over the diversity of browsers truly had to work hard.

HTML5 rings in a new era

This changed fundamentally with the introduction of HTML5. In addition to unambiguous semantics and uniform representation, HTML5 also provides a very comprehensive and standardized JavaScript API. This reduced the browser-specific differences on the one hand and, on the other hand, a few lines of JavaScript were suddenly enough to skip ahead or back in videos, draw diagrams, store data locally in the browser using WebStorage, access the camera of a mobile device or use it to determine the geographical location. This paved the way for Rich Internet Applications (RIAs) with comprehensive, purely client-side functions. The browser mutated from a "visualization vehicle" to a comprehensive runtime environment for JavaScript-based applications in this process.

HTML5 OPENED UP ENTIRELY NEW PERSPECTIVES FOR DEVELOPERS AS WELL AS FOR ATTACKERS.

The diverse functions of HTML5 opened up entirely new perspectives for application developers – but unfortunately the same applies to attackers. Even though security aspects were taken into account in the standardization of HTML5: The scope of functionality in the browser alone means that implementation errors result in security vulnerabilities that are relatively easy to exploit. A relatively quiet paradigm shift is currently underway as well: The application functionality is increasingly shifting to the client, which is precisely the weakest link in the entire security chain.

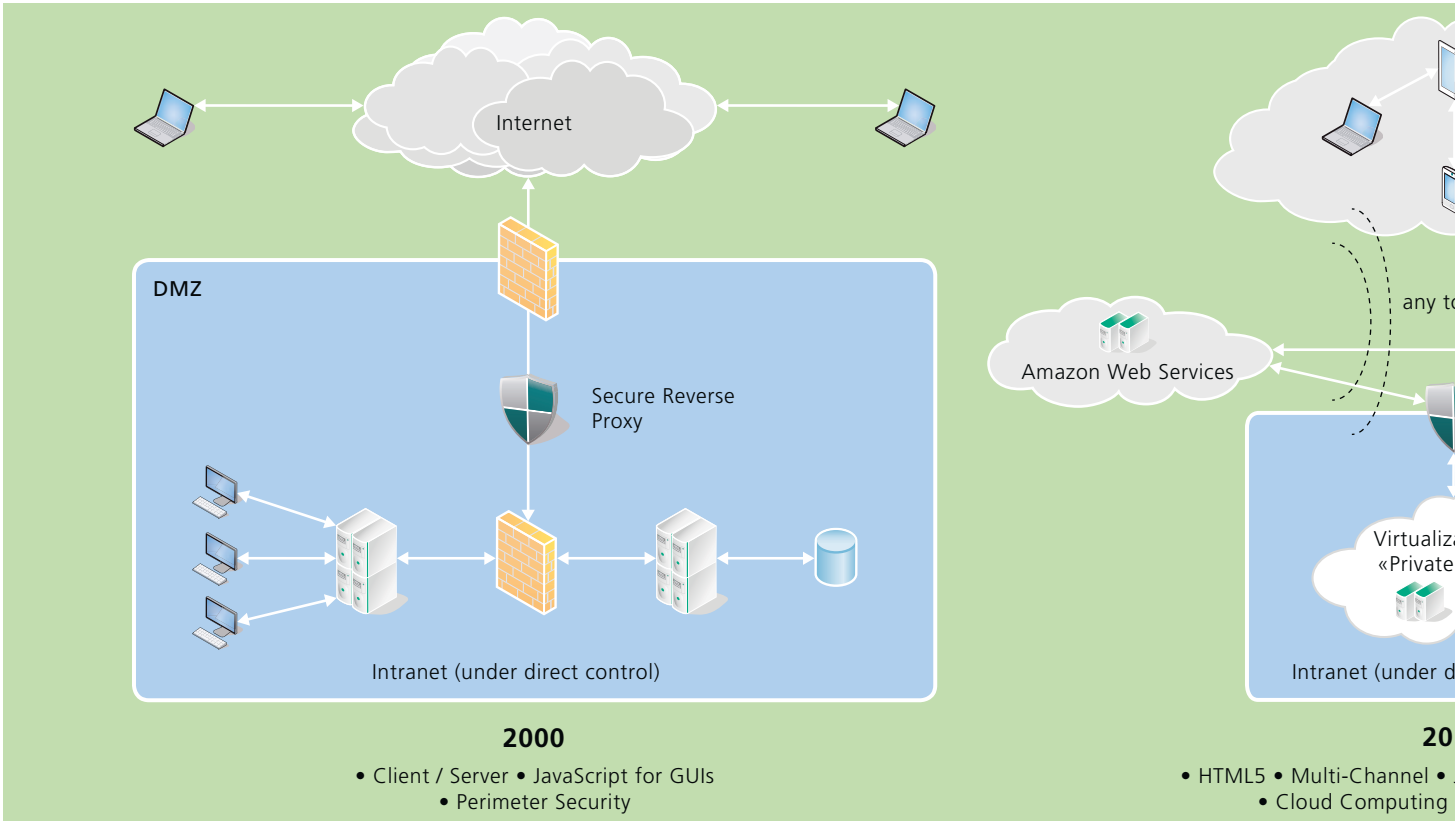

Quo vadis, Internet: Due to the rapid technological development, security needs to be questioned constantly.

ONCE EXTERNAL CODE HAS BEEN SMUGGLED IN, MANY SECURITY MECHANISMS FAIL.

Because of this shift to the weakest link, some established server-side security mechanisms are no longer effective so that protection is reduced or – in extreme cases – eliminated entirely. This is due to the functionality of modern JavaScript frameworks: Due to the client-side logic, server-side security components have only a limited view of the content effectively displayed by the application.

When an attacker succeeds in smuggling its own JavaScript code into these processes, thereby gaining control, things get challenging: Once external code has been smuggled in, many of the internal browser security mechanisms fail since the malicious JavaScript fragments are viewed as a regular part of the application. This makes the comprehensive HTML5 JavaScript API available to the attacker, for example making it possible to operate the camera, access the HTML5 WebStorage content or use Web Workers (JavaScript background processes) to launch a DoS attack.

Security is more central than ever

The examples emphasize that the consistent protection of web applications is more important today than ever before. However, new concepts are required in order to maintain the past level of security. An integrated security concept has to adequately take the client side into account. In addition to clear coding directives for developers, this also includes the consistent use of the existing security mechanisms offered by modern browsers. While these can be centrally activated and controlled by server-side security components (such as nevisProxy), enforcing the security policies remains on the client side. This means a Security Officer has to trust that the browsers apply the specified policies correctly. It is just as or actually more important to consistently apply the protective mechanisms against injection attacks available on the server. Even though these mechanisms no longer offer the same protection in certain scenarios as they did in the past, they constitute an obstacle in the attempt to smuggle malicious code into an application which is not to be underestimated. All security measures jointly pursue one goal: Preventing the infiltration of malicious JavaScript code under all circumstances.

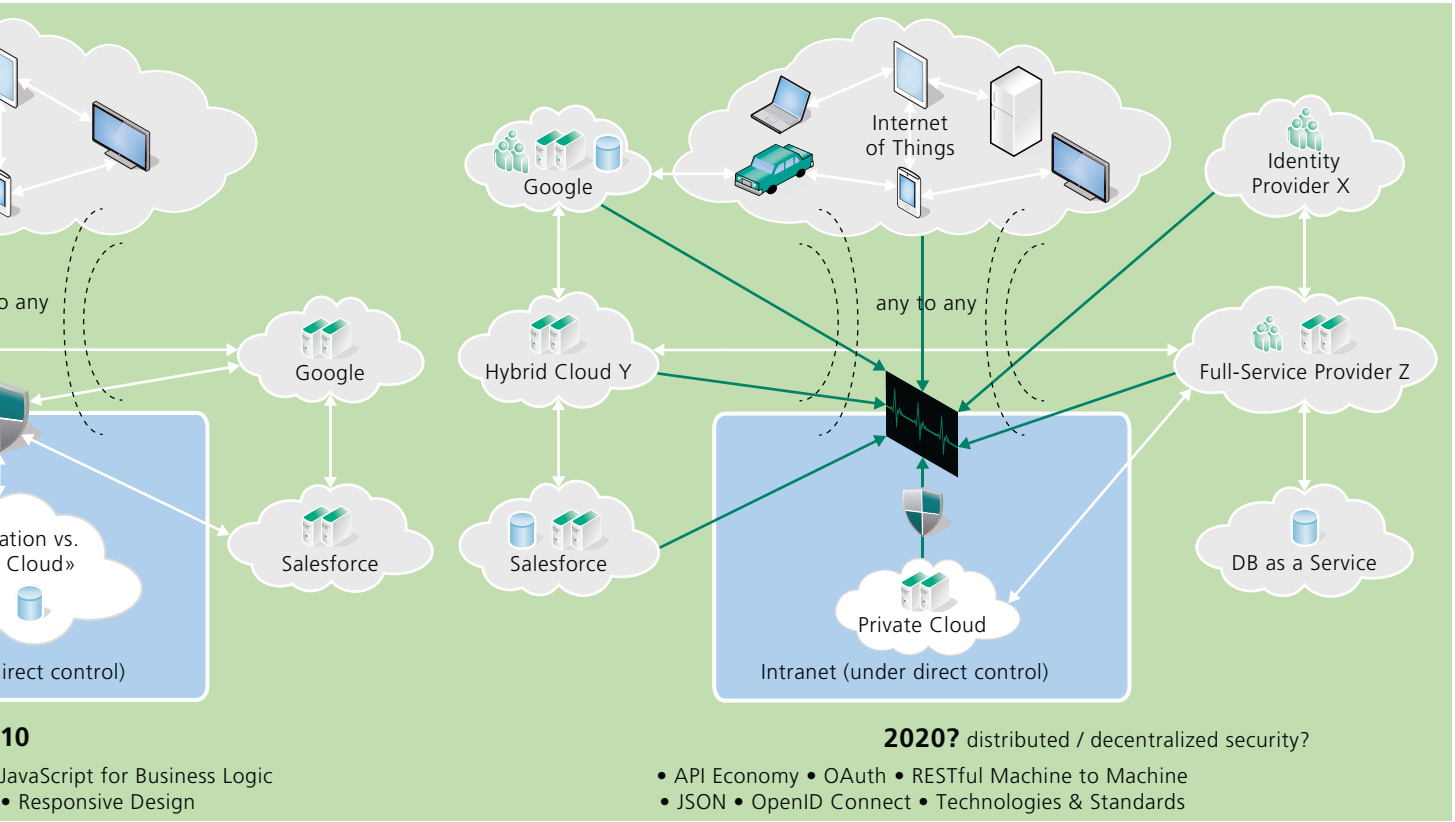

Stephan Schweizer and Thomas Zweifel: Developing strategies for dealing with new Internet hazards.

New security concepts are in demand

Yet one central question remains unanswered by all the preceding measures: How can an organization prepare for the worst case, that is the success of an attacker in the attempt to smuggle its own code into the application? New solutions definitely need to be developed for this scenario. It can be assumed that an infected JavaScript application will behave differently than the original – which is why server-side anomaly detection is a promising approach. The fundamental idea is that there is server-side intelligence which is able to learn and model the normal behavior of an application. Based on this assumption, it is possible to identify deviating behavior and respond to that. The responses to such a situation depend on the respective context. Conceivable options include a message to a central monitoring system, the revocation of critical rights during the session (blocking critical transactions for example), re-authentication of the user or even the immediate termination of the session.

The technical basis for such a system already exists; approaches and algorithms from the field of big data make a valuable contribution here. With a central security architecture based on an entry-gateway approach, the requirements are also met for responding centrally to anomalies that are detected.

The new tasks of the entry gateways

The networking of applications across organizations is usually realized with APIs today. These are web-based interfaces used to exchange data structures in the XML or JSON format. JSON stands for JavaScript Object Notation. As indicated by the name, this is a standard developed for the exchange of data structures in the JavaScript environment. The good news is that networked backend applications and JavaScript-based Rich Internet Applications use the same data exchange technologies. That means the gate-

SERVER-SIDE INTELLIGENCE CAN LEARN THE NORMAL BEHAVIOR OF AN APPLICATION.

way infrastructures can be used for both application types in principle. This reduces costs on the one hand and simplifies operation on the other.

In case of critical or confidential data, access to the APIs naturally needs to be protected by authentication corresponding to the required security level. As applications become more and more closely integrated and linked, but are not necessarily located in the same data center or on the same server, the individual applications must be able to access the respective other APIs in the name of the end user for the purpose of data exchange. The end user should be asked for consent to this procedure; this is known as user consent. Various processes for this purpose are defined by the authentication protocol OAuth. While the identification or authentication of a user by means of an authorization server is part of these processes as well, it is not standardized itself – in part to enable a choice of authentication mechanisms. After authentication, the authorization server issues the access tokens which ultimately enable the applications to access the API.

THE END USER SHOULD BE ASKED FOR CONSENT, WHEN INDIVIDUAL APPLICATIONS ACCESS OTHER APIS.

The functionality to issue the access token and its verification are ideally installed in a level upstream from the applications. Such an architecture has various advantages:

1. Central: Support for the OAuth protocol does not have to be implemented individually for each application. This saves time and reduces costs as well as making the integration of legacy applications possible in the first place.

2. Decoupled: A secure implementation of the OAuth standard imposes high standards on the administration of the issued tokens. Delivering this functionality in an upstream gateway infrastructure enables the straightforward and secure integration of the existing applications.

3. Modular: The authorization server is designed as a central component used by all integrated applications. This component must be able to integrate all existing user directories while continuing to support the existing authentication methods. A central security infrastructure ensures that the required directories only have to be linked once.
4. Additional functionality: Central access token verification is the ideal basis for central accounting and therefore also for new billing and business models.

The existing web entry gateways and perimeter architectures constitute an ideal basis for utilizing these advantages. We are convinced that today's web entry solutions need to be further developed in the direction of API management. In addition to support for the modern federation protocols (OAuth and OpenID Connect), this also includes additional functionality such as accounting, throttling, efficient filtering of JSON content and the selective restriction of the externally published API functions depending on user permissions.

What does this trend mean for Nevis customers?

We consider incorporating new trends a challenge and an incentive. We want to consistently provide our customers with the technical means required for the optimum support of their business in a dynamic environment. That is why the topics and trends identified here have already been integrated into new product functionality or are being considered in future product development.

Designing the new functions focuses on the gradual migration to new application architectures and the parallel operation of the existing solutions – which means that using Nevis ensures maximum investment security and a solid basis to face future challenges.

(For further information, visit the redesigned Nevis website: https://www.nevis-security.ch)

Stephan Schweizer

Stephan Schweizer joined AdNovum as Nevis Product Manager in 2009 and has been responsible for the Nevis product strategy and distribution ever since. He likes to spend his free time in nature with running shoes on his feet and, as a tennis beginner, working on improving his still somewhat shaky serve.

Thomas Zweifel

Thomas Zweifel, MSc in Computer Science ETH and MAS ETH MTEC, has been with AdNovum since April of 2005. As Senior IT Consultant, he advises customers in matters of IT security and IT strategy as well as managing a team of engineers and consultants. He likes to spend his free time on and under the water.

Gion Sialm of the FOITT and Ivan Buetler of Compass Security share their experiences of web security with AdNovum CTO Tom Sprenger.

Browsing the web safely

The security situation on the worldwide web is always in flux. Gion Sialm of the FOITT, Marc Condrau of Health Info Net and Ivan Buetler of Compass Security discuss the current developments with our CTO Tom Sprenger at the security round table.

Attacks are generally known as the greatest adversary to security. Which attacks do the most damage nowadays?

IB: Attacks by Trojans pose the main danger. Companies with large research departments and whose long-term success is dependent on the patents submitted have to deal with constant hackers' attempts to access their intellectual property. Espionage between states is done primarily using Trojan software as well. The second danger is posed by insider stories. Traditional phishing attacks on the web are declining, as they are often difficult and complex. Today, hackers make their way into the Content Management Systems (CMS) of publicly listed companies to obtain their quarterly results hours before they are published. If the hacker team has the know-how required to interpret the quarterly statements, they can do all kinds of highly speculative deals on the stock exchange.

MC: Phishing attacks and Trojans were and still are the most relevant dangers from HIN's point of view. We believe protecting the endpoint, typically the user PC, is crucial. User-friendliness is important here. Users will bypass security solutions with inadequate usability. Another central issue is making the user more aware.

What do today's users expect from an online portal in terms of security?

IB: To me personally, protecting the integrity of my data is important. When using free offers such as Facebook, Twitter etc., I know that I am giving away control and my data is the price. However, I can expect a business application not to pass on my data, to guarantee confidentiality and to stick to the existing regulatory requirements. The standards of the PCI (Pay-

(Marc Condrau's participation took written form.)

ment Card Industry) have to be adhered to when paying online with credit cards. Providers cannot save e.g. the CVC (three-digit code).

MC: Data protection is hugely important in the health sector. Maximum confidentiality and securely identifying communicators is absolutely necessary when accessing and exchanging data. On the other hand, the patient's comprehensive health data has to be available at all times. This balancing act often poses a challenge when designing security systems for the health sector.

Data is being used increasingly across organizations and business systems. How do you secure access within the company?

GS: The (Swiss) federation has clear regulations not just for access management, but also in terms of data protection. It is not permissible to use access data to draw conclusions about a person. If someone e.g. has access to VAT, he probably has his own company. That allows other conclusions to be drawn. That's why we need pronounced data separation. Each office sees only what it manages itself. We secure the data separation not only through regulations and processes but also in technical terms. That makes the solution more complex, but it is important for data protection.

IB: As a rule of thumb, I recommend controlling access near the data, if possible. Programming the web application professionally and safely is also crucial. However, implementation mistakes cannot be ruled out fully in practice. Thus, protecting web applications with an upstream web application firewall (WAF) to prevent such mistakes from triggering a serious accident or super-GAU is a wise strategy. That applies especially when you do not have any access to the web application's source code.

MC: Reliable authentication and authorization of the user is among the other prerequisites for safe, external access to internal date and functions. This is where HIN sets in as an Identity and Access Service Provider for the health sector. The use of validated HIN identities relieves application providers in inter-institutional data exchange. The HIN entry server infrastructure safely establishes the perimeter security of connected applications, e.g. for radiology systems in hospitals, and takes over the application authorization of external users.

PROTECTING WEB APPLICATIONS WITH AN UPSTREAM WEB APPLICATION FIREWALL IS A WISE STRATEGY. (IVAN BUETLER)

If you look at the developments in web applications architecture, a trend towards "app in a browser" stands out. GUIs are no longer sent back and forth in every call, but the whole application runs in the browser (e.g. HTML5). The communication with the backend is turning into technical data communication. How does this alter the demands on a WAF solution?

IB: The technical calls follow standards such as RESTful services, JSON or XML. The WAFs have to be able to understand and validate such protocols and technologies. Problems are occurring with "tunneled" protocols and protocols without open standards (such as ICA or GWT earlier) because the WAF cannot validate them. As long as the web permits proprietary formats for transmitting data, a WAF cannot offer any extensive protection on that level.

Are the new HTML5 applications an issue for the federation?

GS: Yes. We use HTML5 applications in combination with WAF and reverse proxy. We gain flexibility thereby and can intervene in various places for security reasons. Security is crucial to us and thus our supported security features must always be on the most up-to-date level. We are constantly developing our own security architecture and looking for ways to reuse parts of it and what has to be supported anew.

Let's return to data security. There are increasing encounters with application cases in which data leave the protected systems as they are given to third parties. Cloud services or collaborative services used among partners are just some of the examples. The classic perimeter security no longer applies here. Yet, we want to retain control of the data. What kind of approaches are there?

IB: Let's take e.g. e-banking which also presents "foreign" data in the shape of credit card bills. The ownership of the credit card bills does not lie with the banks, but with credit card companies. The e-banking gets a feed from the credit card producer and generates added value for e-banking customers. I see two security aspects here. Firstly, the provider's data (credit card data) can have damaging components. Thus, the bank should validate these data even if they come from a trustworthy partner. Otherwise there is the risk of a so-called second-order injection. Secondly, the credit card company reveals the ownership of the credit card data. What the bank actually does with that depends on the terms of the contract. The customer no longer knows where his data is.

Are there tools for circumventing the conflict of objective between data privacy and data exchange? In terms of government and e-health solutions, aggregated user data would add value. But to do so, people would have to accept their data leaving an organization's sphere of influence.

GS: An aggregation could certainly add value. This is done simply as such datasets need a legal foundation. Thus, if data is aggregated within an authority or across authorities, this is always based on an existing legal foundation. This means that data cannot be analyzed or aggregated without political consent.

> We can achieve a high level of security and flexibility with HTML5 applications and a second line of defense. (Gion Sialm)

MC: In the health sector, the aggregation of particularly sensitive health data with other data is a contentious issue. The federation's e-health strategy takes this into account by using its own patient identification for the patient's electronic file and not the AHV13 number.

Does that mean that companies and organizations have to consolidate their identities and means of authentication in time-consuming and expensive endeavors?

GS: Not necessarily. We have set up a federated architecture and managed to have PKIs, Kerberos, name/password, name/password/SMS and SuisseID authenticated in the federal administration in less than a year.

Do you believe it would be sensible to have an identity pool, which is managed centrally?

MC: Yes, we are fully convinced. HIN offers identities in the cloud as a service for the health sector. If institutions exchange data between each other, the application provider does not have to go to great effort identifying and registering the identities. The user gets access to over 50 applications via single sign-on. We have developed the HIN Access Gateway to facilitate a federated approach.

IB: I believe federated services have great potential in future. Companies will start making their users' identities available on the Internet via services such as Active Directory Federation Services (ADFS) so that cloud providers can profit from these identities. That saves time and money.

An interesting approach. Companies manage the identities and sign contracts with cloud providers. If a new identity is registered, it is automatically provisioned for the use of cloud providers.

IB: And vice-versa. If an employee leaves a company, his account is blocked on the company's own Active Directory and Federation Service. The block takes effect immediately and directly. The user cannot use either the company network or the federated services once his account has been blocked. Unfortuna-

tely, such federated systems are complex. The spread and use needs time to mature.

Federated services have great potential in future. (Ivan Buetler)

MC: Provisioning can also occur via the cloud in the internal IAM. This approach is interesting especially in the health sector where there is high staff mobility. Validated identity data, including information on medical and expert qualifications, can be taken over when a person joins an organization.

This is where the term "dynaxity" is apt – we increase both the dynamics by using federated services and the complexity by aiming for more services. That requires stricter governance...

GS: Definitely. Governance used to be faster than the technology. The monolithic architectures in particular were so complex that we spent years building them up. Today, governance lags behind technology. And the customers are hungry and make great demands. We need a good instinct for what is feasible without making things too complex.

Let's turn to authentication. That is currently an issue in the mobile sector in particular as demands on user-friendliness are generally high. What trends do you notice?

MC: Users of HIN services mainly use mTAN rather than card-based procedures on mobile gadgets. We are currently examining mobile ID.

IB: The authentication has to be above all simple. Complicated, certificate-based solutions such as SuisseID will not take hold across the range in my opinion. The opportunities for such systems lie in closed environments (e.g. the federation) or in applications with higher security needs (B2B or similar).

GS: E-government needs two things: A flexible IAM architecture and a good, simple means of authentication. We have PKI, but that device is as big as the mobile itself. What we like about SuisseID is that the management effort can be outsourced.

What could simple authentication look like?

IB: The user name/password authentication is successful because it's simple. It does not require software. You don't have to buy or install anything and anyone can do it. However, this kind of authentication is not very safe in my opinion. Analyzing the conduct and properties of the client computer when authenticating can increase the security. A so-called "client correlator" examines the settings of the (user's) client computer, e.g. display resolution, installed plugins, browser language, average log-in time when authenticating and the provider's IP range. Such information reveals quickly who was at the computer before the user logs in with his password. In a project named "Panopticlick", the Electronic Frontier Foundation (EFF) has created a prototype, which demonstrates this technology.

MC: A simple user name/password authentication is not acceptable for external access in the health sector. To simplify the procedure for the provider of the application, the HIN platform offers different authentication procedures: certificate-based with Soft Token, mTAN, FMH's Health Professional Card and SuisseID.

Let's spin this thought further. Would solutions that do not require explicit logins be imaginable as long as you're not doing anything risky?

IB: An adaptive security system? That would be a good idea and very comfortable for the user. But in cases of e.g. a VAT repayment, there would have to be a step-up to a higher level of security. And the effort would have to be justified. If the user eventually has to install a client certificate for authentication, he could do that from the start.

If the client correlator gathers the attributes, the assessment is not as sharp. Can the risk be justified in your opinion?

MC: Access without explicit authentication is not imaginable for access to personal data in the health sector.

GS: The federation definitely needs a rethink. In certain cases, such as the VAT repayments mentioned above, such technology can only be used in combination with a procedure that rules out any blurs.

IB: Ultimately, it means that an analysis of data from the client computer can allow predictions about the user. Systems that use such techniques basically know before authentication who has been working at the other end of the line. Such profiling

> PROFILING IS IMAGINABLE TO UNDERMINE ACCESS IN THE EVENT OF UNUSUAL USER BEHAVIOR. (MARC CONDRAU)

systems are already standard in the credit card industry. You pay a few centimes into a pool for damages with every credit card transaction. The credit card companies know precisely when and where a credit card is used. If a debt is incurred in South America while the owner is in Europe, that is automatically recognized and the owner goes unscathed. The credit card system thrives on the latent risk of abuse. This paradigm can be applied to other business cases.

Could the next security generation include profiling, if need be?

GS: As already mentioned, profiling needs a legal basis. Although Swiss citizens increasingly give their personal data to companies or social networks for profiling, it is very difficult to imagine this being allowed in an administration environment.

MC: Profiling in the health sector is imaginable, if access is undermined in the event of unusual user behavior, e.g. when health data is accessed from abroad. The usefulness is highly dependent on the specific application case.

IB: I see great future potential in profiling, i.e. in the compilation and use of users' profiles as well as applications, network traffic and similar. The correlation of data is central to research and security.

As we can see, the possibilities of using data on the web are by no means exhausted. At the same time, users are becoming more demanding and attackers more professional. Thus, protecting our data will remain on the agenda for the foreseeable future. Thank you for this fascinating talk. ■

Round table participants:

Gion Sialm

As head of IAM at the FOITT, Gion Sialm is responsible for access to federal applications. This central service at federation level functions almost exclusively as a trust broker. Internally and externally hosted applications are connected with the biggest authentication directories within and beyond the federation in a flexible manner. Access management remains under the control and responsibility of the commissioning departments themselves.

Marc Condrau

Marc Condrau is a solutions architect and project manager at Health Info Net AG (HIN). HIN was set up in 1996 on the initiative of the Swiss Medical Association (FMH) and the doctors' health insurance fund (Ärztekasse) with the aim of allowing Swiss health experts to use the Internet safely. HIN's core service is identity providing for care providers (approx. 17,000 presently). Secure e-mail services and access control services are offered based on HIN identities. Nearly all established care providers and over 420 institutions use HIN e-mail and over 50 application providers use the access control services.

Ivan Buetler

Ivan Buetler is the co-founder and CEO of Compass Security. Founded in 1999, the company has offices in Jona, Bern and Berlin and employs 35 people. It specializes in ethical hacking, penetration testing and security reviews. Ivan Buetler is an assistant lecturer at the University of Applied Sciences Rapperswil and at the University of Applied Sciences Lucerne. He organizes the European Cyber Security Challenge for Swiss Cyber Storm. He is the intellectual brain behind the Hacking Lab – an international laboratory for security professionals.

Security management in software production

Well-thought-out security management in software production generates genuine added value – both for the customer and in development.

by Marcel Vinzens and René Rehmann

They are called FREAK, Shellshock, Poodle, Heartbleed and BEAST – serious security vulnerabilities in familiar standards and products, which have made headlines regularly in recent years. The National Institute of Standards and Technology (NIST) has recorded over 68,000 software defects known as Common Vulnerabilities and Exposures (CVEs) of various severity in its database. The reported number of these vulnerabilities has increased from around 1000 to 5000 – 8000 annually since the year 2000.

So what do these vulnerabilities and numbers mean for a company such as AdNovum, which produces custom software and software products and frequently also uses open-source software (OSS) as well as licensed closed-source components for its solutions?

The preceding figures and facts clearly show that developing secure software requires a significant effort. In principle, all participants in the Software Development Lifecycle (SDLC) need to assume shared responsibility for the development of secure software. Therefore, it is crucial to approach each phase of the SDLC with the right security mindset. This also includes implementing quality assurance measures and security mechanisms for design, development and distribution as well as monitoring in order to minimize the likelihood of exposure and the effects in case of an exploit. In the following, we show how security management is implemented in our software production which processes more than 100 customer projects annually with over 2000 software deliveries to customers.

Secure software development as the basis

Software security needs to be incorporated in the design and development phases as a matter of principle. Numerous measures contribute to this at AdNovum: Architecture and security sign-offs, quality assurance in the form of manual code reviews and automatic code analyses during the development phase along with training, directives and best practices for developers. In particular, these measures are intended to prevent a situation where security needs to be «integrated» after the fact because the developers focused exclusively on features and usability during software development. Our Security Engineering team is responsible for the control and further development of the measures. It is supported by other engineering teams and continuously takes feedback from software production into account.

AdNovum examines new components for their added value and risk.

As a central element of secure software development, AdNovum operates strict technology and central dependency management as well as using managed repositories for storing software artifacts. What does this mean in concrete terms?

Before AdNovum uses new software components in projects, a critical evaluation of the added value and risk is performed in the course of what are known as technical investigations that form part of the technology approval process. We carefully select third-party components (OSS and proprietary libraries) and store them in centrally managed repositories – provided the evaluation result is positive. All components for every customer solution and product are subject to technology management. We also define a lifecycle status for each component of every version at the operational level. Here we primarily differentiate between components that have been released for use («approved»), are being examined («investigating»), will be replaced («deprecated») and those that are no longer approved for use («forbidden»). Which components are used with what version is reviewed in the course of application development. For components with the status «forbidden» or «restricted»/«investigating» that are not approved for a project, the developer immediately receives an error message. This tells the developer at a glance which components may no longer be used, for example because of

Marcel Vinzens and René Rehmann: Monitoring and assuring the security of AdNovum software.

a known vulnerability. Information specifying the newer version to be used which has been freed of the vulnerability is also provided.

Each internal component, every customer solution and every product is built every night by what is called our NightlyBuild. For projects in wait mode, for example customer solutions with no active further development, we thereby review which components and versions are being used (dependency management). Naturally, the NightlyBuild also performs all tests defined for a customer project or product automatically. We therefore know at all times when a project is no longer current and secure because of lifecycle status changes in third-party components.

ADNOVUM ASSUMES RESPONSIBILITY FOR THE MAINTENANCE OF INTERMEDIATE PRODUCTS AND LIBRARIES ON BEHALF OF ITS CUSTOMERS.

With today's customer projects and products, the objective for reasons of efficiency is to always use existing intermediate products and libraries when possible and sensible, and to «only» develop those components in-house that are either very specific (technical functionality) or result in a competitive advantage for products. AdNovum assumes responsibility for the maintenance of these intermediate products and libraries on behalf of the customer. This also includes for example that we inform customers of the OSS we use in a customer solution. We also notify customers when AdNovum software or OSS used in the same is affected by a security vulnerability. Our security management is responsible for monitoring these components and launching an alerting process as needed. This alerting process regulates how vulnerabilities that are discovered need to be handled and ensures compliance with internal directives.

AN ALERTING PROCESS REGULATES HOW VULNERABILITIES ARE HANDLED.

Active vulnerability monitoring and alerting

Since AdNovum software as described above builds on existing intermediate products and libraries wherever this is possible and sensible, in particular OSS, it is important that any corrections of vulnerabilities in these components are incorporated into the applications delivered by us as well. This is the responsibility of the AdNovum Security Officer, the Security Engineering team and other engineering and product teams.

To identify vulnerabilities, we perform daily active monitoring of the available and known mailing lists and security alerts. When a security alert is issued for an OS component that is used (such as OpenSSL or Apache HTTP Server), or an OS version for appliances, the responsible engineers in a first step analyze whether AdNovum software is directly or indirectly affected by this vulnerability. They estimate the risk and criticality of the vulnerability based on this analysis. If AdNovum software is actually affected and the provider of the affected component has already issued a patch or new release to correct the vulnerability, we provide it immediately or for the next release. If the criticality is high and no fix is available yet, we develop our own or at least find a workaround so the vulnerability cannot be exploited. We can for example deactivate specific elements of the affected components that are not necessarily needed by the AdNovum software. The

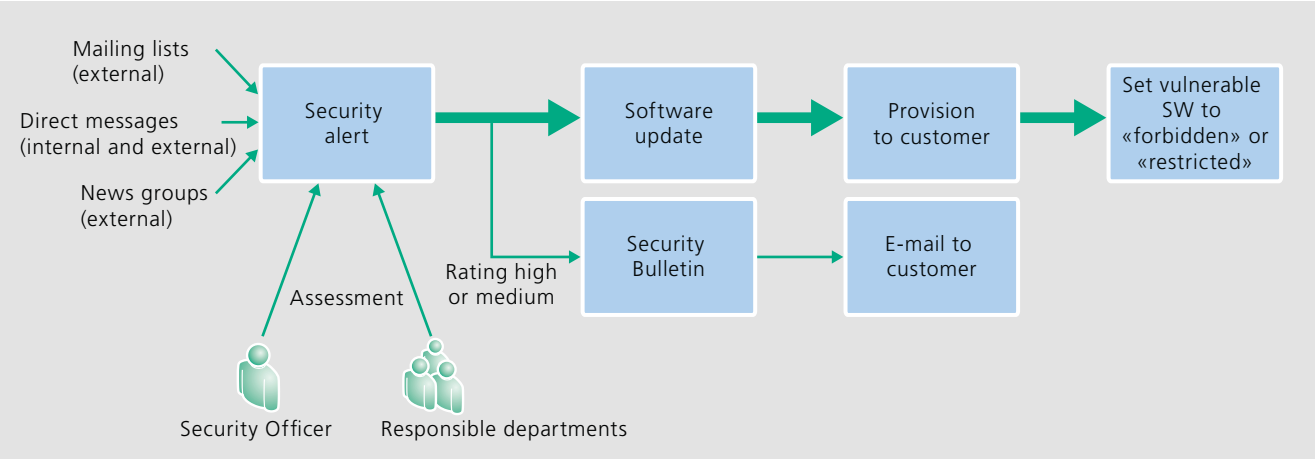

Whenever vulnerabilities arise, AdNovum follows an established process.

new release is automatically tested in the NightlyBuild or in a continuous/daily build depending on the criticality. With supplementary manual tests we ensure that the new release is no longer affected by the vulnerability so that the quality expectations of our customers are met. It is especially important to ensure that the fix is not limited to the vulnerability as such but also eliminates its cause.

For security alerts of high or moderate criticality, we send a Security Bulletin to our customers. It includes our risk assessment along with recommendations for the update priorities. We also provide the release dates for the individual affected software products.

In case of highly exposed components such as OpenSSL or Apache HTTP Server that are often used in the access zone, we send out a Security Bulletin even in case of lower or no criticality.

THE FIX MUST NOT LIMIT ITSELF TO THE VULNERABILITY, BUT ALSO HAS TO ELIMINATE ITS CAUSE.

Customers often monitor the more common OSS products themselves using the same sources. Questions on the evaluation of criticality often come up. That is why our Security Officer coordinates communication with customers to ensure consistency. The respective project managers at AdNovum are responsible for project-specific communication.

Once the new releases of the affected components have been delivered, we internally set the components containing the vulnerability to «forbidden». This ensures that these components are not unintentionally used again in any customer project or product. In exceptional cases we set the affected versions of the components for certain projects and products to «approved restricted» when it is clear that they are not affected by the vulnerability directly or indirectly. This may be the case when the defective element of the affected component is not used in these projects and products, and if it is not opportune from a customer perspective to update the component at the time. Customers with components of AdNovum's Security Suite Nevis are not only informed with the Security Bulletin in case of vulnerabilities, but also through the Nevis blog.

Conclusion

While software that is free of defects and security vulnerabilities is something to strive for, it cannot be guaranteed in the real world. It needs to be the goal of any reputable software provider to minimize these vulnerabilities and defects through the principles, tools and processes of secure software development. A software company also needs to know on the one hand which components are installed in what customer solutions and, on the other hand, which weaknesses currently affect these components.

When security problems arise or become known, professional conduct and clearly defined directives for handling vulnerabilities offer the best protection. That means companies need to clarify who is responsible for security problems, for example a Security Officer. This is especially important for software companies. Security monitoring and a defined, established alerting process with consistent external communication are needed as well. In order to accomplish this, the processes for preparing new releases of customer projects and products as well as the (automatic) testing of these releases have to work effectively and efficiently. Products such as the Nevis Suite on the one hand and all customer projects on the other hand benefit from these core elements of security management at AdNovum. That is why security management in software production not only means costs for customers with a software maintenance agreement, but also genuine added value. ■

René Rehmann

René Rehmann, Dr. phil. nat., has been working for AdNovum as Business Project Manager and Security Officer since 2012. He focuses on security and identity management in projects. In his free time he likes to frequent the golf courses of this world.

Marcel Vinzens

Marcel Vinzens, with AdNovum since 2002, holds an MSc ETH in Computer Science and is CISSP-certified. As Technical Project Manager and Solution Architect, he was dedicated to the conceptual design, engineering and integration of security and middleware systems for several years. He has been monitoring the lifecycle of the products and frameworks used by AdNovum internally and in customer projects as Deputy CTO since 2013, advises and supports customers in security and architecture matters and has been heading the Security Engineering team since the beginning of 2015. He spends his free time with his family, enjoying culinary highlights and leisure activities.

It depends on the context: How good does authentication need to be?

Not only people but also apps and things access information today. That makes the issue of security more important than ever.

by Martin Kuppinger

About KuppingerCole

KuppingerCole, founded in 2004, is a global analyst firm focusing on Information Security and Identity & Access Management (IAM). Governance, Risk Management and Compliance (GRC) is another core area of KuppingerCole research. KuppingerCole as an independent analyst group organizes conferences, seminars, workshops and webcasts in the fields of information security, IAM and GRC. It also hosts the European Identity & Cloud Conference, which has established itself as the main event for opinion leadership and best practices for Identity & Access Management, Cloud and Digital Risk in Europe.

The IT security directives of many companies state that content with a higher classification such as «confidential» may only be used with two-factor authentication. But in a world where more and more users are accessing applications, systems and therefore content with a wide variety of devices, such requirements have long since become inadequate.

Security is no longer purely an issue internal to companies

The traditional focus of information security was on employees accessing internal systems. Yet the scenario has changed fundamentally in the last few years. Applications no longer run just in the internal data center but also in the cloud. More and more applications are not only being opened up to business partners but to customers as well. Long since has access not been just from desktop PCs but from numerous devices at many different locations, often over WLANs accessible to the public. Apps working over what are known as APIs (Application Programming Interfaces) are increasingly being used as well. These APIs are gaining importance for integration into the Internet of Things (IoT) and for networking business processes between companies.

> More and more applications are not only being opened up to business partners but to customers as well.

In other words: There is more and more access as time passes, not only by people but also by things, apps and other systems. Information needs to be protected, even in this complex IT reality. Access needs to be adequately authenticated and authorized. Authentication has to determine whether the communication partner is the person, system or thing it purports to be. Authori-

zation is about the decision what system and information access to grant in concrete terms.

It is apparent that the basic differentiation between less critical with simple authentication, typically with a username and password, and critical with two-factor authentication is no longer sufficient in this environment. Things can and have to authenticate themselves differently than people, for whom in turn the device they are using is a factor in determining how authentication can actually be realized.

The balance of two aspects is decisive

How strong or weak authentication should be depends on the balance between two aspects. One aspect is the context in which access takes place: What device is being used? Where is it being used? Is the installed anti-malware current? What networks are being used? Are there any indications of abuse? We are talking about the risk associated with access. This risk may be greater or lesser, depending on the context.

SO THERE IS NO RIGHT OR WRONG, BUT ONLY APPROPRIATE AUTHENTICATION.

The other aspect is the risk that is acceptable for an interaction or transaction. This requires a differentiated understanding of risk, which goes beyond the classification of documents, data or applications.

Furthermore, very different methods have to be supported depending on the access channel. When a user accesses a cloud service through an app, different methods and standards apply than when an employee accesses internal business applications from a notebook.

Authentication has to take the context into account. There is no right or wrong, strong or weak authentication, there is only appropriate authentication. Common two-factor authentication with a Smartcard or OTP Token may be insufficient for highly critical transactions.

High time to rethink the approach

Companies need to adapt their strategies, rules and implementation for authentication and authorization to the changed reality. Directives have to become more flexible and must be based on concrete risk. Practical methods to determine risk are needed. Various authentication methods have to be enabled to provide flexible support for different devices. Authorization needs to become variable. More or less may be permitted depending on authentication.

This requires new directives and concepts, but also flexible applications that integrate different authentication methods and standards in addition to supporting a broad range of use cases. It also means that applications must be able to work in a context and make their authorization decisions based on the authentication strength as well.

IT IS HIGH TIME TO RETHINK AND MODERNIZE APPROACHES FOR THE OPTIMUM PROTECTION OF INFORMATION.

It is high time to rethink and modernize existing approaches for the optimum protection of information. This applies all the more since the risk of attacks has increased massively in recent years and no end to this trend is in sight. Where this journey to adaptive authentication and authorization under consideration of context information will lead was a topic avidly discussed at the European Identity Conference 2015 held in Munich, Germany this May (cf. www.id-conf.com).

Imprint

Publisher:
AdNovum Informatik AG
Corporate Communication
Röntgenstrasse 22
8005 Zürich
Phone +41 44 272 6111
E-Mail info@adnovum.ch
www.adnovum.ch
Responsibility and editing:
Andrea Duttwiler
Feedback: notitia@adnovum.ch
Design and realization:
Comuniq, Zürich
Photography:
Gerry Nitsch, Zürich
Printed on Balance Pure