



PRIVACY BY DESIGN  
@ Västtrafik



**OWASP**  
The Open Web Application Security Project  
2014-04-24

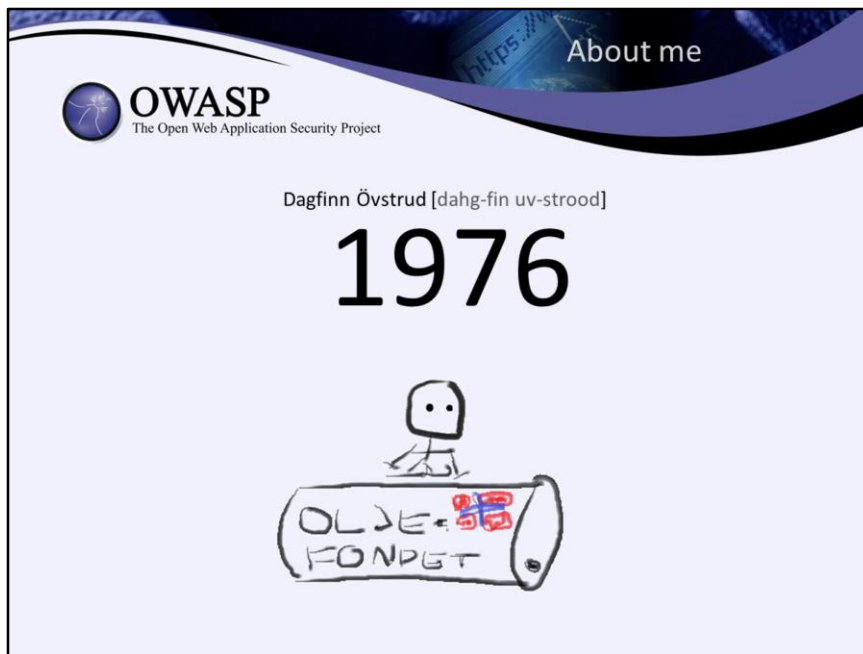




# CONTENTS

- Introduction (whoami/Västtrafik)
- Systems in this presentation
- Protection
- Challenges
- Final thoughts

These will be the topics of this talk. I'll give a short description of the systems I have selected, so you get some sort of a grasp on what they are and what they do. Then I will describe some methods of protection used and some of the challenges faced in this regard, and at the end I will give you some final thoughts on the subject of privacy and security.



Hi!

My name is Dagfinn Övstrud, and I am an IT architect who used to work at Västtrafik. And Västtrafik will be the subject of this talk, and more specifically, how some of the systems used by Västtrafik handles and protects privacy data.

I was born in Oslo, Norway in the year 1976. As all Norwegians, I was born with a pair of cross-country skis on my feet, and was wrapped in my first lusekofte shortly after birth. And of course, I expect the Norwegian oil fund to provide for me when I retire. As you will later learn, I worked only for nine years in Norway, but that should still be long enough to have me covered pretty good.



I started my career at only 17 years of age, at a company called Akers Mic (retailer of home entertainment, music etc). Not as an IT employee, but this is where I discovered the joy of computers, and luckily I had a cool (and veeeery brave) boss, and by 1999 I had climbed from the bottom up to being the IT Manager.



I met my Swedish wife in Oslo back in 1995, and at that time I was (aside from being a workaholic) in a metal band called Neurosilence. I made (a rather stupid) promise to her; That if my band didn't have a record deal by the year 2000 I would move with her to Sweden. And obviously, the record deal never happened. The summer of 2002 we moved to Västra Frölunda, and a year after that to Hönö, where she grew up.

I stopped playing guitar completely, but last year I picked it up again. The result of that can be found on <https://www.facebook.com/CircleOfIndifference>



In 2003 I started working at Banverket as a combined secondline resource and a MS SQL "expert". After working my way up the ranks for a few years, I moved over to Västtrafik in 2008. What is it with me and public transport?



In 2014 I turned myself over to the dark side, when I joined Kentor as a consultant. My main area is Microsoft products (MS SQL Server, IIS, TMG etc), but I also have some experience with stuff like Varnish, Squid, Cisco ACE, two-factor auth etc etc.



I have a special interest in security, high availability and performance. And I am not afraid to speak up either. As you can probably guess, that makes me less than popular at times, in the eyes of project managers, developers and others. But I stand my ground!






Guru



Me



A very non-scientific approximation, so don't waste any time and energy analyzing this further. It may just as well be:





**OWASP**  
The Open Web Application Security Project

About



## Responsible for public transport in Västra Götaland

[vest-rah yuh-tah-lahnd]

## Wholly owned by VGR

[veh-geh-err]

Until a couple of years ago, Västtrafik had 50 owners. These were the 49 municipalities in Västra Götaland, and the region itself. Now, VGR is the sole owner.



**OWASP**  
The Open Web Application Security Project

About



## Board of directors:

Leif Blomqvist (S), Västra Frölunda, ordförande  
Mimmi von Troil (M), Göteborg, vice ordförande  
Tore Hult (S), Alingsås  
Britt-Marie Andrén Karlsson (S), Ellös  
Ulf Olsson (S), Borås  
Conny Johansson (S), Falköping  
Jan-Erik Wallin (M), Vara  
Eva Abrahamsson (M), Smögen  
Maj Steen (M), Borås  
Soili Brunberg (MP), Hisings Kärra  
Max Andersson (MP), Göteborg  
Nanna Siewerts Tulinius (FP), Lerum  
Ewa Hamberg (V), Göteborg  
Elving Andersson (C), Uddevalla  
Benny Strandberg (KD), Kungälv

Want change in the public transport  
sector? Talk to *them!*



**OWASP**  
The Open Web Application Security Project

About



Coordinator, responsible for timetables,  
routes and journeyplanning, and more.  
Lots more

All traffic is procured and contracted

Västtrafik does NOT operate traffic directly, it is contracted and run by third parties.



**OWASP**  
The Open Web Application Security Project

About



Operations are 50% tax-funded, the remainder is ticket-funded

**Non-profit!**

What does this mean? Well, when ticket prices are raised, it is NOT because the VP wants a bigger bonus, it's because running public transport isn't free. It also means that any profit were to be made, it is returned to the owners. That means YOU (indirectly).



**2.700** busses, trains, trams and ferries

**22.000** stop points

**283.000.000** journeys made during 2013



A typical weekday:

**390.000** kilometers driven (roughly 9,6 times around the Earth)

**75.000** visitors on [www.vasttrafik.se](http://www.vasttrafik.se)

**1.500.000** request sent to the travel planner

9,6 times around the earth. In one day. That is alot, people. Alot.



- Events in the ticketing system
- Customer information and actions related to travel cards (registration is voluntary)
- Customer service/CR-tickets
- ...And a bunch of other stuff not related to the systems in this particular presentation






**OWASP**  
The Open Web Application Security Project

# SYSTEMS IN THIS PRESENTATION

- The ticketing system (VIX)
- VTK [veh-teh-kaw]
- [www.vasttrafik.se](http://www.vasttrafik.se) / Mina Sidor [mee-nah see-duhr]



To be fair, many issues are more the result of the pricing model than the ticketing system itself.

The developer and supplier of the ticketing system used by Västtrafik is VIX/ERG, a company hailing from Perth, Australia. Staff from VIX sat in-house at Västtrafik until recently, whereupon they moved out to offices close by. The system is located in-house, but everything above the hardware-level (apart from backup and anti-virus etc.) is operated by staff from VIX.

The ticketing system is a contact-less card system based on the MIFARE-chip. There is no real-time communication between the validator devices and the back office system, events are uploaded in batches and saved to a proprietary database.

Travel card events are identified by card number (also printed on the physical card, like a credit card).

It has absolutely no access to, or interest in, customer information (it is card centric), and by itself completely anonymous for all intents and purposes.

Relevant data is copied to interested parties in separate, non-proprietary databases.




## Linking customers to travel cards



VTK is a directory, connecting customer information to a travel card number and events in the ticketing system. Or so it was in the beginning. It has since grown to be a lot more than that. When you register your travel card, this is where your personal information is stored.



## GUI for Customer Relations



Has its own graphical user interface used by Customer Relations. It's not a CR system, but rather an interface where CR can look up a customer and events pertaining to a customer's travel card.



Exposes a number of web services used by [www.vasttrafik.se](http://www.vasttrafik.se) and a few other systems. These web services contain methods for logging in, reviewing and updating customer information in VTK, manage the customers travel card and fetch travel card history (journeys etc.), and more. These web services are NOT publicly available, they are used by other systems only.

In addition to this, it sends SMS and mail to customers (confirmation of payment, payment reminders, ticket information etc.)

Handles autoladda (automatic reload of period card or card balance) and other card related actions in communication with internal and external systems.



vasttrafik.se is built on the EPI Server platform, running on IIS on Windows. In your face, Heartbleed!

“My pages” creates a user account in VTK when a customer registers his/her first travel card. A customer cannot be created without an association to a travel card. But, once created, the card can be deleted from VTK without affecting the user account. This is a perfect example of a system who’s original scope was very different from current use. Company accounts are handled slightly differently, they can register without a card, but workarounds are implemented to overcome the card requirement.

Integrations include VTK, maps, travel planner, traffic information (both traffic disturbances and planned traffic date) and more.

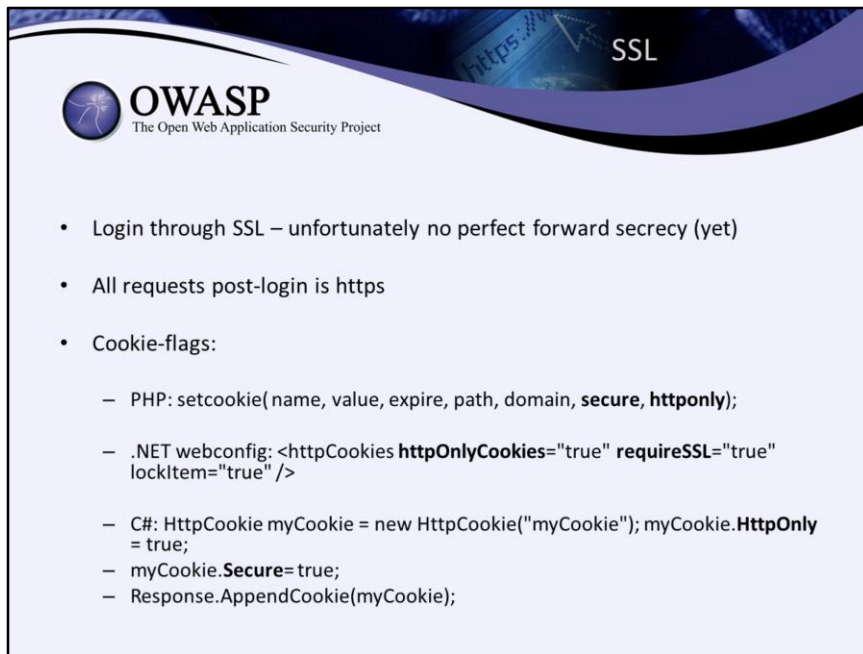
On “My pages” you can for example:

- Register and block travel cards
- Administer “autoladda”
- Generate travel history reports
- Check travel card balance



# PROTECTION

- SSL
- Network segmentation and firewall rules
- Authentication/access control
- Logging/Auditing
- Culling/purging data
- "To cloud or not to cloud"

A presentation slide with a blue and white wavy header. The header contains the OWASP logo (a globe) and the text "OWASP The Open Web Application Security Project" on the left, and "SSL" on the right. The main content area is white and contains a bulleted list of points regarding SSL and cookies.

**OWASP**  
The Open Web Application Security Project

SSL

- Login through SSL – unfortunately no perfect forward secrecy (yet)
- All requests post-login is https
- Cookie-flags:
  - PHP: `setcookie( name, value, expire, path, domain, secure, httponly);`
  - .NET webconfig: `<httpCookies httpOnlyCookies="true" requireSSL="true" lockItem="true" />`
  - C#: `HttpCookie myCookie = new HttpCookie("myCookie"); myCookie.HttpOnly = true;`
  - `myCookie.Secure= true;`
  - `Response.AppendCookie(myCookie);`

The login page is https, as is the posting of the login data. All subsequent requests, as a logged in user, are carried strictly over https. One could choose to always use https, regardless of login status, as today this doesn't incur that much of a relative penalty in performance. But, as is the case with Västtrafik, if you run a separate cache layer in front of the web servers, and this happens to be something like Varnish, you would have a problem. Varnish does not support https, meaning you would have to offload SSL in front of Varnish. The hardware handling SSL offloading (a hardware load balancer) did not support setting or clearing the secure cookie flag, which would be required in order to ensure that the cookie is securely transferred between the client browser and the offloader, and still readable by the webfront over http. It might support it now, or maybe it even did back then, but it's difficult for me to test properly without access to it. In addition to this, we did not know how much extra stress this would put on the loadbalancer during extreme peaks. And much of the ssl-traffic IS cached anyway, it's just the requests that talk to serverfarms that need the cookie that bypass Varnish. Besides, it's a surefire way to avoid caching private content by mistake.


Maybe next time, all this will be taken into consideration from the start of the project. We (or they) are still learning!



Secure: The cookie is sent over https only

HTTPOnly: Protects the cookie from being read by a client-side


LockItem: When included is the web.config in the root of the site, no other web.config or application configuration can override these settings.



**OWASP**  
The Open Web Application Security Project

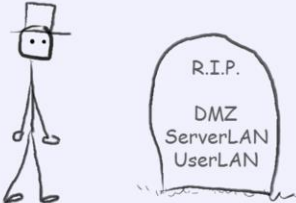
## Network segmentation and firewall rules



- Default "deny any any"
- Only necessary ports, protocols and directions are opened between subnets
- A firewall between subnets adds the ability to control, monitor and audit network traffic



This is an example of logical segmentation and system separation. Each bubble represents a subnet, in the case of the reverse proxy, load balancer and the internet-facing firewall the bubble represents two or more subnets (outside/inside etc.).

The Layer 3 switch you see at the top is maybe not really relevant in this context, but I included it anyway, as a small tip on how you can protect your infrastructure against certain types of DDoS-attacks (syn-floods etc) and other unwanted traffic. You can use it to NULL-route (also known as "blackhole filtering") that traffic before it affects your infrastructure. The only (still bad, though) effect of the attack is potentially a loss of bandwidth for external access. But nothing on the inside is affected, which means zero effort and time to get the systems back online when the attack is over, in other words you reduce or completely remove any collateral damage.

Instead of NULL-routing it you can send it to a separate log server with TCDdump for analysis and forensic purposes (connected directly to the switch and completely separated from the rest of the infrastructure).

The old design practice of setting up a DMZ, a server LAN and a client LAN should be avoided when possible. A much safer design is to separate every system by network segmentation (sub netting), and use in/out access lists in a firewall that has "deny any any" by default, and any access to any other system must be opened by request, when such communication is necessary.

Know your protocols, ports and in which direction your network traffic travels! I have lost count over how many times we've had developers and system suppliers facepalm themselves when they've been informed of these design principles, and accuse us of delaying implementations because they need to more or less reverse engineer their own system to make it work at

Västtrafik. Most of the time we end up searching for denies in the firewall logs to sort it out, but at least we try to have them figure it out by themselves first. They need to learn!

Separating systems this way does not only dramatically reduce the attack surface in your infrastructure, it has the added benefit of giving you the opportunity to gather data for analysis, auditing, real-time monitoring, forensics etc.

Another thing to keep in mind is to be in control over privileges needed for system access, internal system communication, integrations etc., so that you easily can assign the least amount of privileges necessary on all components, from the start. Service accounts seldom need to be local administrators and database owners. And changing this on a system once it's in production is a lot more difficult.




**OWASP**  
The Open Web Application Security Project

#derp

From [http://www.experts-exchange.com/Hardware/Networking\\_Hardware/Switches/Q\\_28348607.html](http://www.experts-exchange.com/Hardware/Networking_Hardware/Switches/Q_28348607.html) regarding the use of subnets:

"If all you have is one site and you know that you will always have less than 254 devices on that network, then sub netting and super netting is probably a waste of time and effort."

"The security argument you heard mostly applies to WAN connections. Since WANs travel over a greater geographical distance, they can sometimes be more susceptible to mischief."



Ah yes... The internet is full of experts that are willing to share their wisdom.



- User accounts in "Mina sidor" are mapped to internal user accounts in VTK (these are completely separate from GUI user accounts, which are Windows user accounts)
- VTK's GUI has strict role-based access control
- VTK's web services require authentication
- No enduser-systems can access the ticketing system directly

User accounts in "My pages" are mapped to internal user accounts in VTK (these are completely separate from GUI user accounts, which are Windows user accounts).

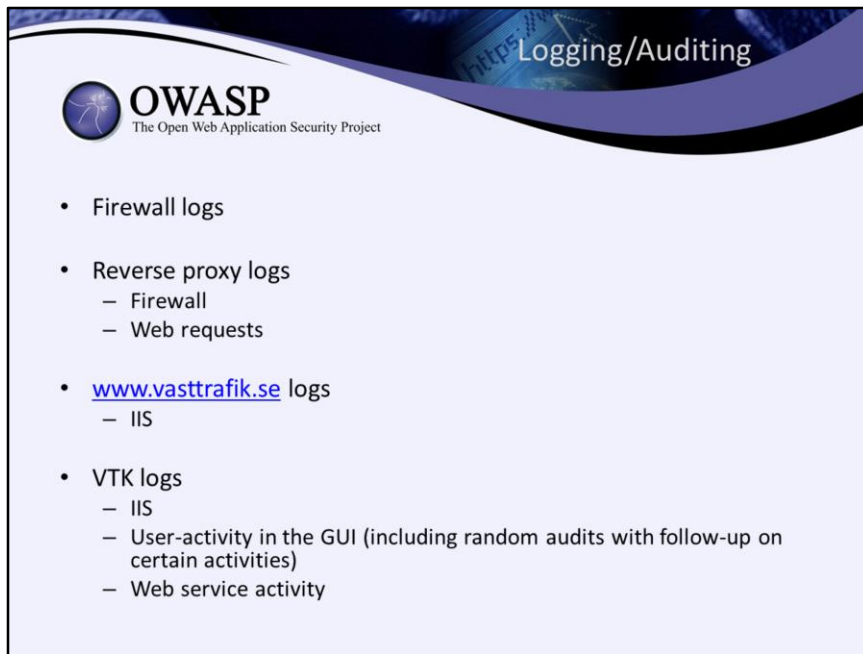
VTK's GUI has strict role-based access control, depending on the function of the operator. Each operator is assigned one or more of these roles, which are based on personal windows accounts and security groups. Each person accessing VTK is by these means given the least amount of privileges necessary to perform their tasks.

CR have strict rules and guidelines regarding lookups and sharing of this information, in accordance with PUL. Lookups are only allowed when requested by the customer. No travel history is given over the phone, but card balance can be if requested. On the other hand, the customer has a right to check if the correct sum has been withdrawn on the card on a particular journey. If travel history beyond that is requested, the customer will either receive it by regular mail to the address registered in VTK, or in person with ID.

If the card is not registered in VTK, the travel history will not be given out under any circumstance. This solution is not perfect, but as close as it can be without being unreasonably impractical. For instance, there is today no way of controlling that the person who travels with a registered card is the same person who registered it.

Even the police and other authorities does not have access to this information without providing a written request, which is then subjected to a confidentiality examination before any information

is handed over.

A presentation slide with a blue and white wavy header. The header contains the OWASP logo (a globe with a star) and the text "OWASP The Open Web Application Security Project" on the left, and "Logging/Auditing" on the right. The main content area is white and contains a bulleted list of log sources.

OWASP  
The Open Web Application Security Project

Logging/Auditing


- Firewall logs
- Reverse proxy logs
  - Firewall
  - Web requests
- [www.vasttrafik.se](http://www.vasttrafik.se) logs
  - IIS
- VTK logs
  - IIS
  - User-activity in the GUI (including random audits with follow-up on certain activities)
  - Web service activity

Firewall logs gives you a nice overview of traffic both in real-time, and historically. It does however, result in massive amounts of data, so for follow-ups or dashboards I would recommend installing a log handling system. SPLUNK is probably one of the best there is, if you can afford it...

Västtrafiks reverse proxy acts as a firewall and a web proxy, and both these functions generate logs. The website itself generates webserver logs, and events within "My pages" is logged to a database on the application level.

VTK has, in addition to server logs, a function that logs actions performed in its GUI (who and what).

Random auditing is performed on a regular basis, and alerts are triggered on any suspicious or abnormal activities. This may be viewed as an invasion of the CR agent's privacy, but it is my belief that this is actually more of a safety for them, as it means that they cannot be accused or suspected of doing anything wrong when they have not done so. Unless they have been careless with their account information or left their screen unlocked when leaving their desk that is.



**OWASP**  
The Open Web Application Security Project

Culling of data

- VTK
  - Travel history and other card activities are purged from VTK's copy of the ticketing systems database based on the following rules:
    - Irrelevant events are deleted immediately
    - Travel history is deleted after 90 days
    - Other card transactions are deleted after 605 days
- Ticketing system
  - No data is purged
  - Card numbers are anonymized after 90 days, so they cannot be traced back to an individual or a physical card

The ticketing system registers a lot of events that are not of any value or interest to VTK. These are deleted from VTK's copy of the ticketing system database on a daily basis, and are never available to anyone.

Travel history, that is card validations performed when entering (and sometimes exiting) vehicles, are kept as-is in 90 days for follow-ups by CR (when a customer case requires it), and for customer reports available through "My pages". Any transactions of this nature older than 90 days are deleted daily. Deleted, NOT flagged as old and still kept in the database. Still available to malicious experts until the record is physically overwritten in the data- and log files, but if that is a problem, you already have other problems.


Other card transactions, such as increasing card balance, has to be kept available a little longer. Currently this period is 605 days (for customer service reasons). When these transactions reach this predetermined age, they are deleted in the same manner as the other transactions.

The ticketing system itself does not delete any data, it has to be available for archive and statistical reasons. But, the card numbers will undergo anonymization after three months, and the routines for this is being discussed and hopefully implemented in the near future. This will be done by using a secret algorithm to remove the actual card number and replace it with




something else. This new number cannot be used to connect historical information in the ticketing system to a customer in VTK. More on this later.

Where is your (i.e. the customers private) data stored?



In which region or country is your data stored? Certain types of data may require storage within the national borders. And how can you be sure that this is the case when it's "The Cloud"?

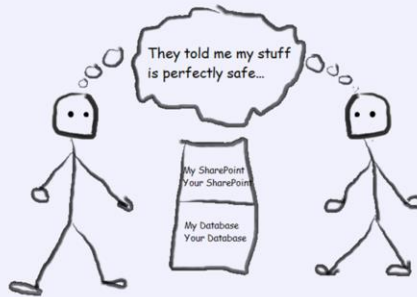
## Who and what has access to it?



Security is somewhat of a catch 22. I know this point can be argued, but bear with me. Without access to detailed information on how they operate and secure their sites in regards to physical access and network access, details regarding their employees, infrastructure, other systems on the same sites etc. you cannot be 100% sure that everything is as they claim. But, if they DO show you all this, their security is by this very action already compromised... Not by much, but still. And they are not likely to share this information either way.

YOU (or your company) are still responsible for protecting the data, even if you outsource operations or cloudsource it.

## Shared infrastructure?



Obviously you will share the infrastructure with the cloud provider's other customers, the question is at what level. Same virtual server? Same Hypervisor? Same network segment?

”But we have signed a Service Level Agreement  
and a Non-Disclosure Agreement!”



SLA's and NDA's are all fine and dandy on management level, but the reality is that when the shit hits the fan, they are seldom worth the ink they are printed with. If (or when) something happens, you may receive a small symbolic compensation, but I strongly doubt that the amount is on the same level as the actual cost for you (which in many cases is difficult or near impossible to estimate). And for the people whose private data has been compromised, any such agreement is totally worthless. What do they get? "Uh, gee, we are so sorry..."

The Cloud is perfect for a lot of things, but not everything. If you handle private data in anyway, especially on a government- or municipal level as a public function, this data should be kept in-house. Or at least in a private cloud on a site within the country, where you have the possibility to perform full-scale audits.

Another problem with cloud-based services is that when a major incident occurs, it's on a scale far bigger than that of your own infrastructure, meaning that the potential downtime is increased, especially if you are on the bottom of the priority list.

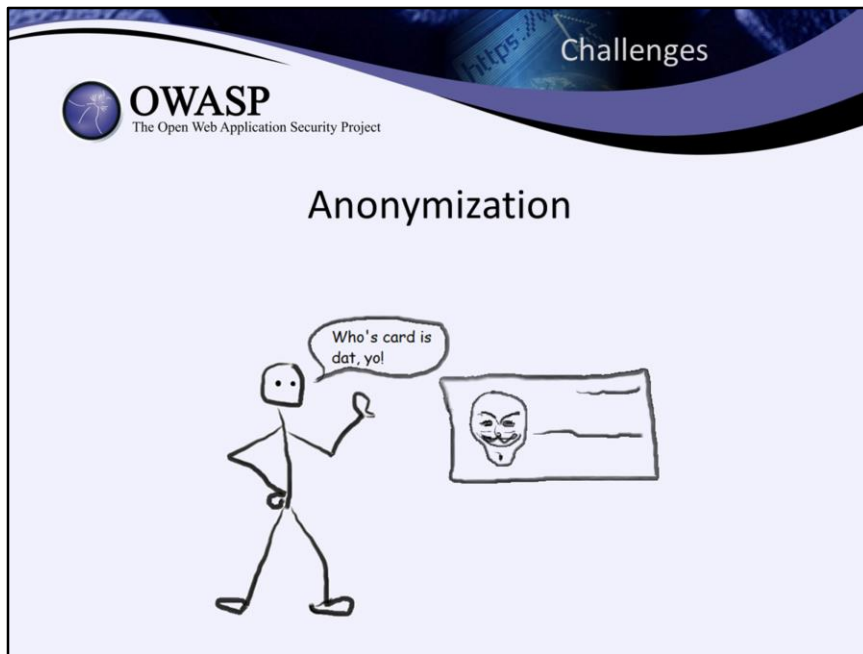
All of the above can be said about regular outsourcing as well.



**OWASP**  
The Open Web Application Security Project

# CHALLENGES

- Anonymization
- Culling/purging data *everywhere*
- Backups
- Security
- Legacy



Completely irreversible anonymization is unfortunately impossible, if you want to maintain the possibility of analyzing historical travel patterns etc. To achieve this, the algorithm responsible for re-writing the card number must make sure that any card always receive the same obfuscated number each time it runs on new sets of data.


One of the reasons one needs to keep the data as complete as possible might be to adhere to the archive act, and in this case any future public research (samhällsforskning).

But what do I mean when I say you can't have the process completely irreversible? Well, the point is that certain individuals will always have access to data prior to *and* after anonymization. This means that they can monitor one or several specific card numbers in fresh datasets, and compare rows before and after anonymization to create their own translation table, and by this method have access to those specific cards complete history. But, the number of individuals with the opportunity to do this is extremely limited, and the risk level is assessed to be adequately low to be acceptable compared to the practical value of maintaining historical data.

When it comes to the process itself, since the data to be anonymized exists in one specific table in a database, the method is pretty straight forward. All one needs to do is to have some sort of protected and encrypted service, SSIS-package or similar, with some sort of one-way rewrite algorithm, and run it on a daily basis anonymizing the tail of the "live" data according to the time period set by the business rules, which in Västtrafik's case is data older than 90 days.



## Culling/Purging (swedish: "Gallring")




One challenge when it comes to culling ticketing system data is that it is copied over to several other databases outside its own boundaries. In VTK's case that is easily solved by business rules deleting data, but other systems are dependent on historical data, and maintain a full copy. For anonymization to be meaningful, any data copied outside the source system must also follow the same anonymization rules and utilize the exact same algorithm. Otherwise, if data needs to be reloaded from the initial source, the anonymized card numbers would differ wildly, and the whole concept would crash and burn.



### What about backups?


Now what!?



What about backups? Given the fact that the anonymization process only affects data a given period back in time, more recent data consists of actual card numbers. These will continuously land themselves on backup devices as-is. It would be extremely difficult, completely unreasonable, very expensive and time-consuming to anonymize backed-up data, so in this case one is allowed to treat backups as offline or inaccessible data that is not subject to the same rules regarding anonymization and purging. In other words, it's ok to pretend it doesn't exist.



## Security



It is extremely important that any system handling data of a private nature is properly updated with the latest security patches, anti virus definition files etc., even more so the, if any, components are exposed to the internet. Even the infrastructure hosting the system plays an important part in this respect, you do not want any adjacent insecure systems to be able to access any other system if compromised.

## Inheriting old crap, what do you *really* need to store



As mentioned end previously, VTK's original scope was far smaller than current use. One aspect is the fact that you need to provide your "personnummer" if you want to register. This was primarily because Västtrafik wanted to sync customer-address with SPAR. If that was really necessary is debatable, but now that this is built in to the system on such a fundamental level, even if it was desirable to remove it, it would be costly and somewhat problematic to do so.



## FINAL THOUGHTS

- Posterboys are rare
- It's not "*all or nothing*"
- Security By Design as a policy
- Spread awareness and knowledge
- EVERYONE needs to contribute!

Västtrafik is by no means the posterboy of privacy by design. The systems I have presented have come about at different points in time, designed and implemented by different companies and different people. The still relatively high standards are not strictly the result of an agreed-upon policy forcing these principles on any system handling privacy data, but rather the personal know-how, integrity and dedication found in the members of the IT department. And in part, pressure from The Swedish Data Inspection Board. Over time, and as we speak, this is slowly settling as standard design principles being written down as a general company policy, and eventually this will be a complete suite of policies to be followed at all time. But, all of it probably not singled out as policies in regards to privacy in particular, but rather design principles to be followed regardless of data, as it really applies to sound security design in general. Privacy data is not the only data you have worth protecting.

You do not have to have everything perfect, anything is better than nothing. Many regard the task of secure design daunting, and feel that since they cannot get it a 100% secure it's not worth doing.

Another important aspect is awareness and knowledge. By that I mean that not everyone has to have a PhD in computer science in order to be of any use in the field of security, but rather have the understanding that there are real threats out there, there will always be threats out there, and anyone can be a victim of these threats, and that tools already exist to aid in these matters. A general misconception is that an attack is directed at a specific target. While that is true in some cases, many incidents are random attacks attacking anyone vulnerable. So the argument "who on earth would want to attack us?" is not really valid.

Like I said, not everyone has to be a genius. It doesn't have to be like climbing Mount Everest or be expensive like a trip to Mars to design and implement systems in a secure manner. *Security* is not scary, but the *lack of it* is. What people need to learn is that we can let the geniuses design

and build the wheels, and that all the rest of us need to do is use these wheels to our advantage. Or even use carts and wagons built by the semi-geniuses' using the geniuses' wheels. There is no shame in that!

The primary mission of OWASP is sharing and spreading knowledge in the field of security, and that extends beyond the contributors and board members. All of YOU, as members of OWASP, have a responsibility as well, to do the same when you talk to your colleagues, business management, customers etc. etc. Raise the general awareness! It's a bit like swimming upstream with bricks tied to your feet, but resigning to ignorance is not the answer.



Thank you!



**OWASP**  
The Open Web Application Security Project

- Contact information:

- [dagfinn.ovstrud@kentor.se](mailto:dagfinn.ovstrud@kentor.se)

- [dagfinn.ovstrud@gmail.com](mailto:dagfinn.ovstrud@gmail.com)

- <https://www.facebook.com/CircleOfIndifference>