Cloud Computing Security

Raining on the Trendy New Parade

Andrew Becherer, Alex Stamos, Nathan Wilcox

BlackHat USA 2009

iSEC PARTNERS

Agenda

- Cloud Computing Defined
- Software as a Service
- Platform as a Service
- Infrastructure as a Service

Special Thanks

- Chris Clark
- Alex Vidergar
- Scott Stender
- Andreas Junestam

Cloud Computing

"am i the only one who has an urge to punch myself in the neck whenever i hear about 'the cloud'?"

- Arshan Dabirsiaghi

Commenter at Jeremiah's Blog

No, Arshan, you are not the only one.

iSEC PARTNERS

Cloud Computing

- Term is useless
- What is it not?
  - Virtualization
  - Remote backup
  - Most of the stuff called cloud computing

Cloud Computing

- Generally means:
  - Lots of general purpose hosts
  - Central management
  - Distributed data storage
  - Ability to move applications from system to system
  - Low-touch provisioning system
  - Soft failover/redundancy

- If you aren't re-writing your software, it's not Cloud Computing

Cloud Computing

- All technological and policy assessments must be based on:
  - Specific deployment model
  - Specific implementation
- Anybody who talks about "Cloud Computing Security" in general is selling you something

Software as a Service

Authentication

Audit

Taking Back Control

Software as a Service

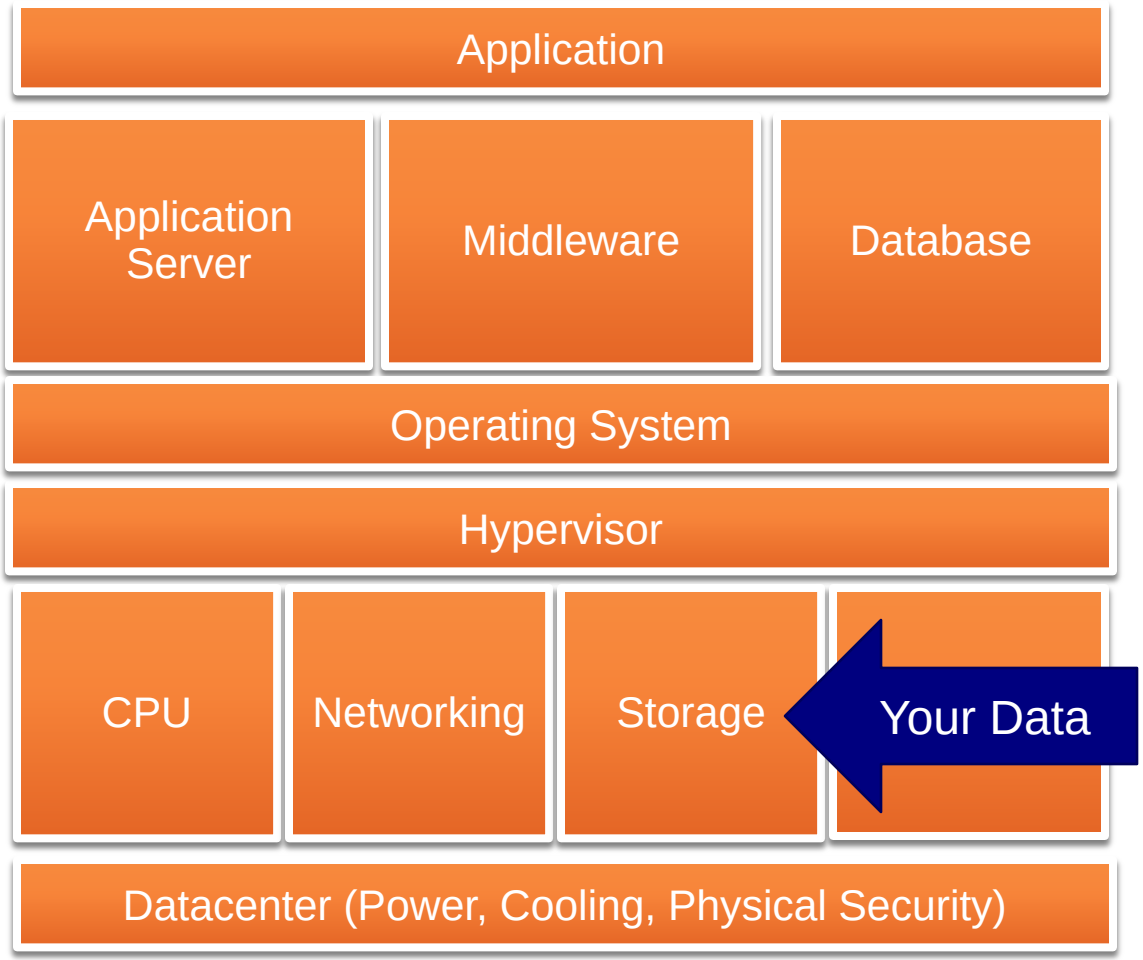

iSEC PARTNERS

Cloudy Authentication

- Recent Twitter incident reinforces an important point:

"*No matter how low an opinion you have of your users, they will figure out a way to disappoint you.*"

-Stamos' Law

Authentication and Credentials

- What controls do we lose when using SaaS?
  - Physical and logical network barriers
  - Endpoint restrictions and management
  - Non-password auth
  - Fine grained credential quality controls
  - Password reset process
  - Real-time anomaly detection

- Most IT departments believe in some of these
  - Many people doubt usefulness of perimeter
  - Hackers aren't unicorns

iSEC PARTNERS

Account Quality

- Some services mix consumer accounts with "datacenter admin"

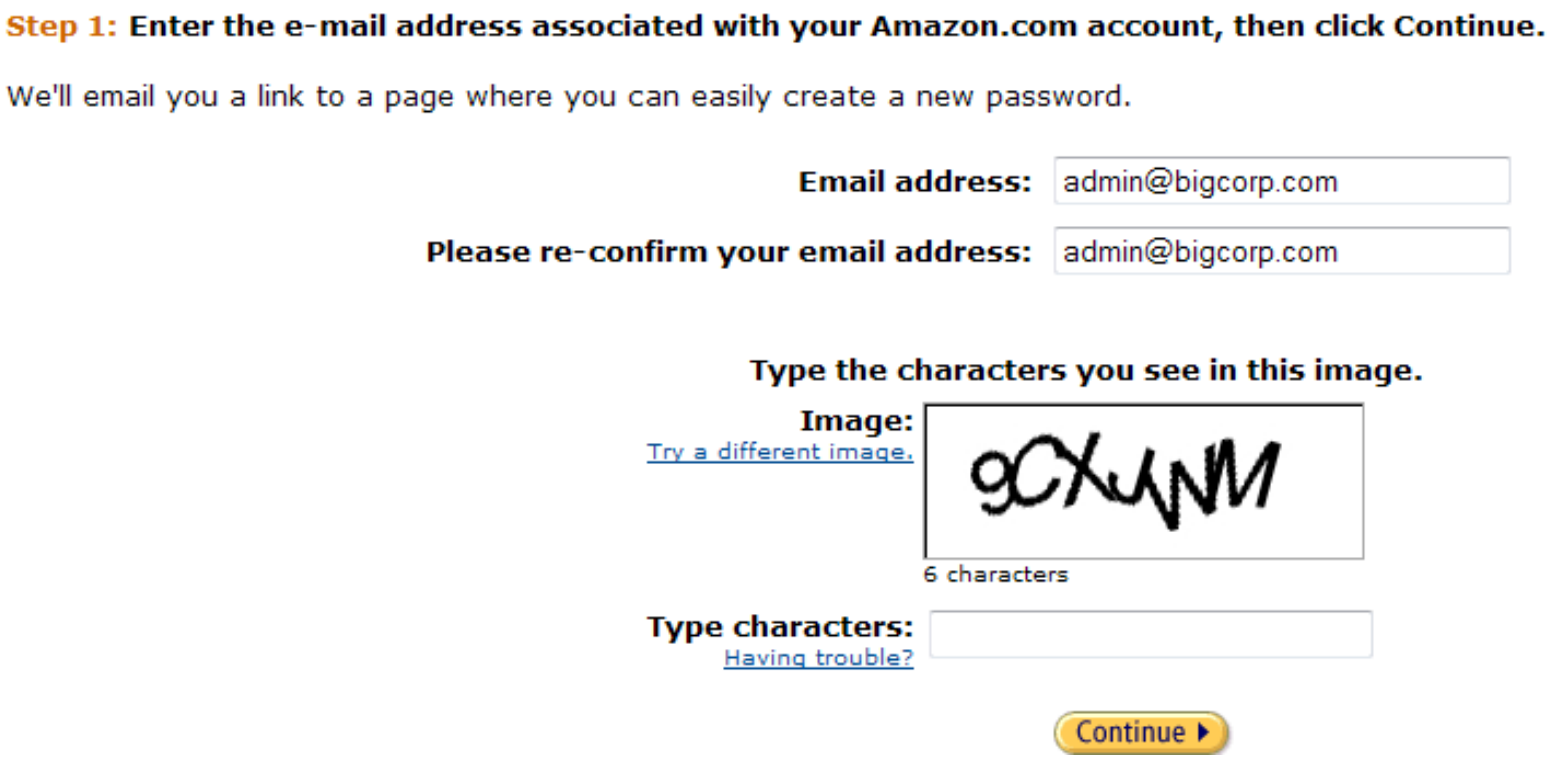

iSEC PARTNERS

Audit and Logging

- Most SaaS vendors do not provide the level of audit logs necessary to recover from a serious breach

- What do I need to know?
  - Who logged in?
    - When?
    - From where?
  - What administrative actions were taken?
  - What documents/data was accessed?

SaaS Audit Comparison

	Login Events	Admin Events	Data Read	Data Write	SSO
Google Apps	**No**	**No**	**No**	**Yes**	**Yes**
Office Live	**No**	**No**	**No**	**Yes**	**No**
Salesforce	**Yes**	**Yes**	**No**	**Yes**	**Yes**

- Missing from all these guys:
  - Per record/document read records
- Salesforce has much more centralized data access

iSEC PARTNERS

Google Apps Audit Logs

- Google provides users with some self-service history:

Access Type [?] (Browser, mobile, POP3, etc.)	**IP address [?]**	**Date/Time** (Displayed in your time zone)
Browser	75.212.205.36 *	11:48 am (0 minutes ago)
IMAP	208.54.5.51	11:26 am (21 minutes ago)
Browser	24.120.153.162 *	8:57 am (2.5 hours ago)
Browser	24.120.153.162 *	8:33 am (3 hours ago)
Browser	24.120.153.162 *	7:52 am (3.5 hours ago)

- Admins can see last logged in time

- Google claims information available via DocList API

iSEC PARTNERS

Salesforce Audit

- SF.com provides detailed login, admin event logs

Source IP	Login Type	Status	Browser	Platform	Application
75.212.167.64	Application	Success	IE 8	WinNT	Browser
75.212.167.64	Application	Failed: Computer activation required	IE 8	WinNT	Browser
75.212.167.64	Application	Success	Chrome 3.0	WinNT	Browser

- Write logging available in Force.com DB, not read

iSEC PARTNERS

Credential Alternatives

- Some providers offer mechanisms to return login control to you
- Google offers SAML integration:

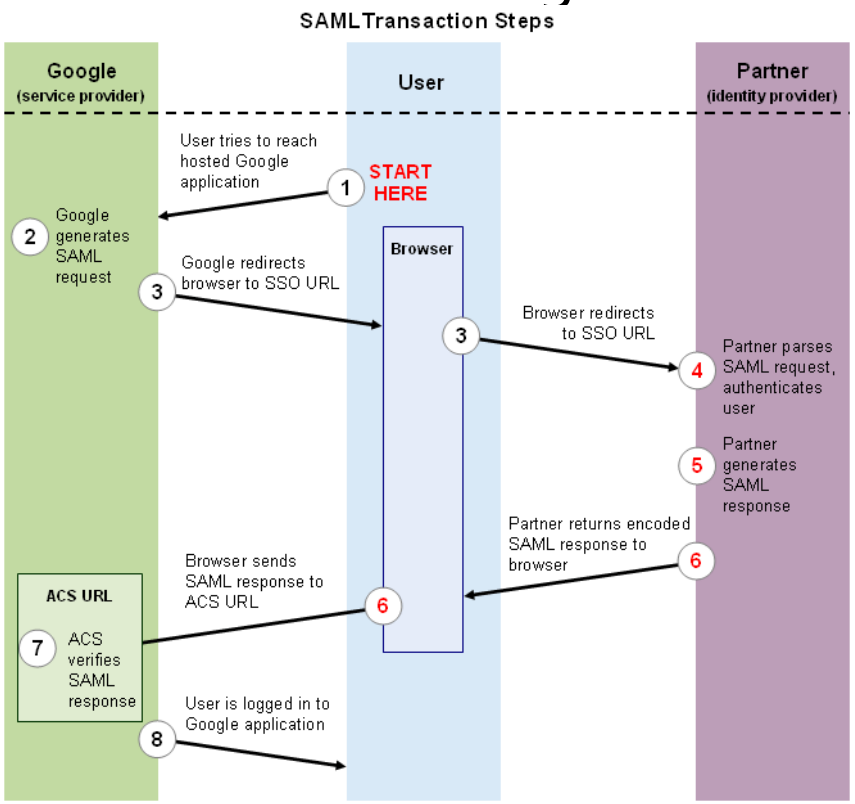

http://code.google.com/apis/apps/sso/saml_reference_implementation.html

Why take back authentication?

- Doesn't it defeat some of the benefits of the cloud?

  Yes.

- But it allows you to:
  - Use alternative cred scheme (token, cert)
  - Completely control password policies
  - Implement internal password reset
  - Perform anomaly detection on login attempts
  - Place the portal behind VPN
    - Access control
    - Endpoint management

iSEC PARTNERS

SaaS Auth Bottom Line

- Recommendations:
  - Strong policies on quality and rotation
  - Employee education is key
    - Never re-use credentials
    - Anti-Phishing techniques
- Use off-site SSO if available
  - Consider additional restrictions using VPN
- Map to what protections you had pre-cloud

iSEC PARTNERS

Legal and Regulatory

- Thank you to:
  - Joe Gratz – Durie Tangri Page Lemley Roberts & Kent LLP
  - Jennifer Granick and Kurt Opsahl – EFF

- IANAL. Any mistakes are my own. Get a good lawyer.

- Read for yourself:

http://aws.amazon.com/agreement/

http://www.google.com/apps/intl/en/terms/user_terms.html

http://www.salesforce.com/company/msa.jsp

iSEC PARTNERS

Legal Concerns: Liability

- As you would expect, Cloud EULAs promise nothing
- What happens in case of…
  - Breach
  - Data loss
  - Disaster
  - Business event
- You can't expect these folks to take on financial liability, but it would be nice if they would promise to help

iSEC PARTNERS

Legal Concerns: Self-Testing

- Most of the EULAs specifically disallow malicious traffic
- Important part of IT security, sometimes required
- Amazon, assured us that they are ok with pen-testing with the owner's permission
- Salesforce, Google allow app-level pen-testing of hosted apps

iSEC PARTNERS

Legal Concern: Search and Seizure

- Does using Cloud Services decrease your protection from search of your data by:
  - Law Enforcement?
  - Civil Plantiffs?

- The answer seems to be **YES.**

iSEC PARTNERS

Legal Concern: Search and Seizure

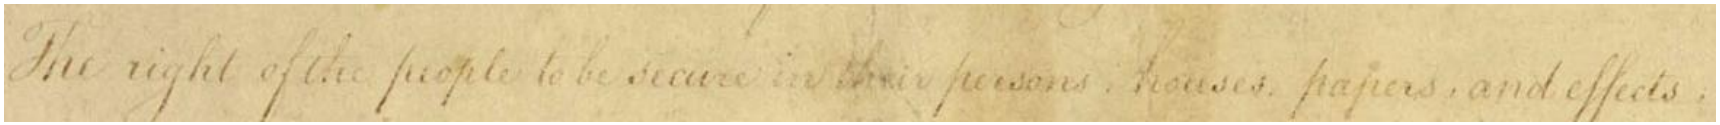

"The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no Warrants shall issue, but upon probable cause, supported by Oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized."

- Apparently "persons, houses, papers, and effects" does not include "hard drives in Google's DC"
- Several statutory protections, but mostly only protect "communications"
  - Are your Salesforce data "communications"?

iSEC PARTNERS

Legal Concern: Search and Seizure

- What do you lose in the Cloud?
  - Protection of a Warrant
    - Signed by Magistrate
    - Requires "probable cause"
  - Guarantee of notice
  - Ability to fight seizure before hand

"*Storing data yourself, on your own computers — without relying on the cloud — is the most legally secure way to handle your private information, generally req uiring a warrant and prior notice. The government asserts that it can subpoena your data from cloud computing providers, with no prior notice to you.*"

-Granick and Opsahl, EFF

More info:
https://ssd.eff.org/3rdparties/protect/storage

iSEC PARTNERS

Google's Response

"Google complies with valid legal process. Google requests that all third-party legal process be directed at the customer, not at Google, and we provide our customers with the tools and/or data required to respond to process directly. If Google directly receives legal process concerning customer or end-user data, it is Google policy to inform the customer of said process, unless legally prevented from doing so. We are committed to protecting user privacy when faced with law enforcement requests. We have a track record of advocating on behalf of user privacy in the face of such requests (including U.S. Dept. of Justice subpoenas). We scrutinize requests carefully to ensure that they adhere to both the letter and the spirit of the law before complying."

Platform as a Service

Developers are the Essential Audience

The Contenders

Attack Surface Case Study

Summary Comparison

Platform as a Service

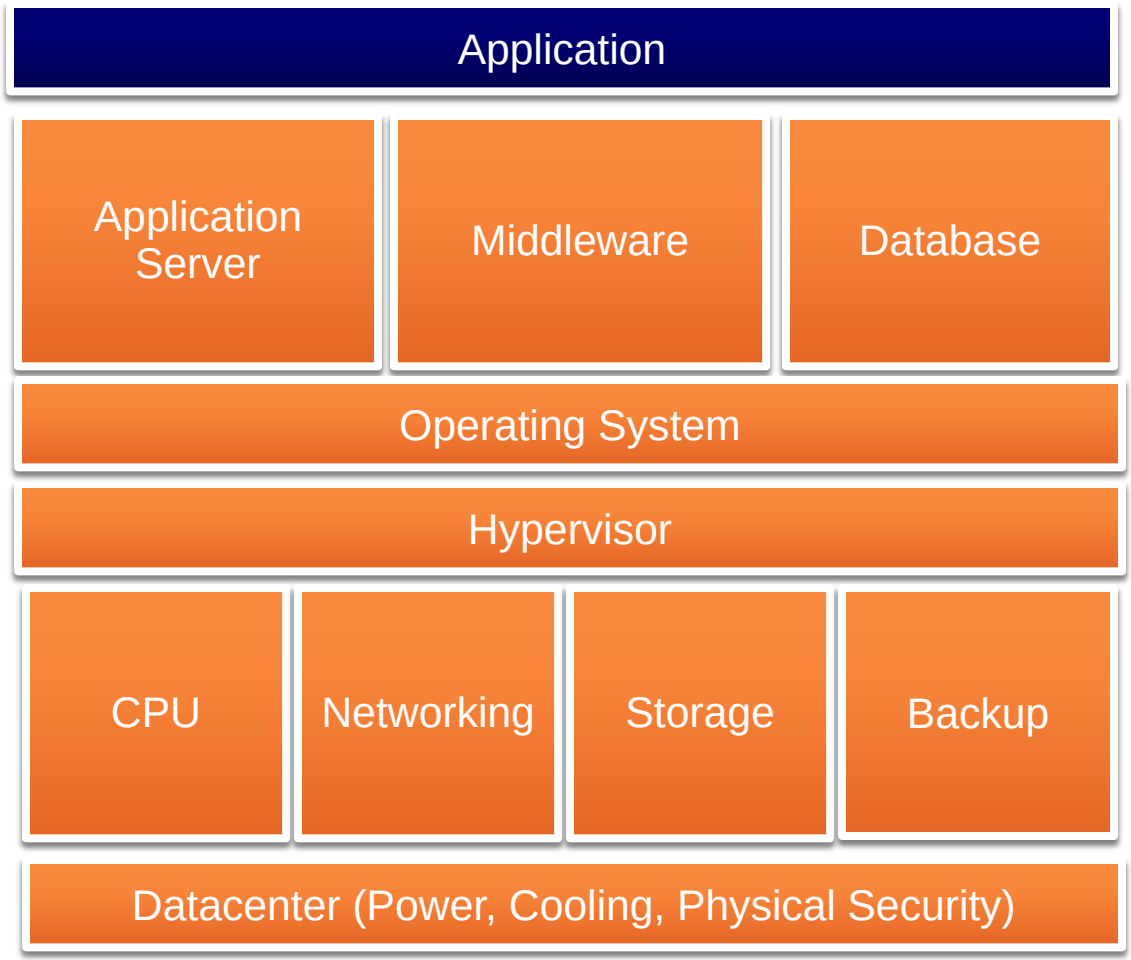

iSEC PARTNERS

The Contenders

Google
AppEngine

Salesforce
force.com
platform as a service

Microsoft
Windows Azure

iSEC PARTNERS

Attack Surface Cases

- Questions to consider:
  - Secure out of the box?
  - Is it {hard / easy} to get {right / wrong}?
  - How could it be better?
- Selected cases:
  - CSRF
  - XSS
  - SQL Injection

iSEC PARTNERS

Cross-Site Request Forgery

- Subtle, often misunderstood.
- Can be mitigated almost transparently.
  - Frameworks can tie forms to sessions.
  - Just remember to confine modifications to POSTs.

GAE CSRF Prevention

- Not easily found in documentation.
- … nor the discussion groups.

- Django mitigates CSRF with configuration.
- App must be configured to use Django in lieu of default framework.

GAE CSRF Prevention…

- Step 1: Switch to Django on AppEngine

```
# Force Django to reload its settings.
from django.conf import settings
settings._target = None

# Must set this env var before importing any part of Django
os.environ['DJANGO_SETTINGS_MODULE'] = 'settings‘
```

iSEC PARTNERS

Django on AppEngine: http://code.google.com/appengine/articles/django.html

GAE CSRF Prevention…

- Step 2: Enable CSRF Protection on Django

```
MIDDLEWARE_CLASSES = (
    'django.middleware.common.CommonMiddleware',
    'django.contrib.sessions.middleware.SessionMiddleware',
    'django.contrib.auth.middleware.AuthenticationMiddleware',
    'django.contrib.csrf.middleware.CsrfMiddleware‘,
)
```

iSEC PARTNERS

Django CSRF protection: http://www.djangoproject.com/documentation/0.96/csrf/

Azure CSRF Prevention

- Azure-specific docs on CSRF not found.
- ASP.NET best practice is per-session ViewState MAC:

```
void Page_Init(object sender, EventArgs e)
{
    ViewStateUserKey = Session.SessionID;
}
```

- Important: Ensure the session is shared across VMs.
  - See the SDK's AspProviderDemo.

iSEC PARTNERS

ASP.NET CSRF docs: http://msdn.microsoft.com/en-us/library/ms972976.aspx#viewstate_topic12
Code and pitfalls: http://www.owasp.org/index.php/.Net_CSRF_Guard

Force.com CSRF Prevention

- All standard controls protect by default.
- Developers may unwittingly miss protection:

```
<apex:page controller="myClass" action="{!init}"></apex:page>
```

Preconditions *exclude* CSRF token validation

Input parameter vuln to CSRF

```
public class myClass {
  public void init() {
    Id id = ApexPages.currentPage().getParameters().get('id');
    Account obj = [select id, Name FROM Account WHERE id = :id];
    delete obj;
    return ;
  }
}
```

Ref: http://wiki.developerforce.com/index.php/Apex_and_Visualforce_Security_Tips

CSRF Lessons

- Deviations from the "ancestor" frameworks lead to configuration headaches:
  - AppEngine/Django middleware.
  - Azure requires new session configuration.
- Force.com trades a better default at the cost of learning a new language and platform.
- Custom handlers tend to inadvertently disable the protection.

Cross-site Scripting

- We focus on output encoding over input validation.
- Requires more developer awareness than CSRF.
  - Typically devs must consider which parameters to escape.
- The framework solution is not inherently different in the Cloud environment (as with CSRF).

GAE XSS Filtering

- AppEngine templates can use Django filters, including an XSS encoder: escape
- This filter encodes for HTML body and non-JS attribute contexts.

iSEC PARTNERS

GAE Templates: http://code.google.com/appengine/docs/python/gettingstarted/templates.html

GAE XSS Filtering…

- Example:

```
<html><body>
   {% for greeting in greetings %}
     <p>Santa says: {{greeting.content|escape}}</p>
   {% endfor %}
<p>Today's limerick is:
<input type="text" value="{{limerick|escape}}">
</p>
</body></html>
```

iSEC PARTNERS

GAE Templates: http://code.google.com/appengine/docs/python/gettingstarted/templates.html

Azure XSS Filtering

- Azure relies on standard ASP.NET for output encoding:
  - HttpUtility.HtmlEncode or .InnerText property.
  - HttpUtility.UrlEncode
- HtmlEncode and .InnerText are for the HTML body or non-JS attribute contexts.
- Examples:
  - Welcome1.InnerText = "Hello, " + User.Identity.Name;
  - Response.Write(HttpUtility.HtmlEncode(Request.Form["name"]));

iSEC PARTNERS

XSS tutorial: http://msdn.microsoft.com/en-us/library/ms998274.aspx

Force.com XSS Filtering

- Two UI frameworks:
  - S-Controls provide UI via JS mechanisms (older design)
  - Visualforce provides markup template language

Security Tips: http://wiki.developerforce.com/index.php/Apex_and_Visualforce_Security_Tips

Force.com S-Controls

- Eschew eval() and string-based callbacks:

```
window.setTimeout('flingAnimal(' + evilParam + ')')
```

- Use higher level DOM api and .innerText, *not* .innerHTML
- Server-side expansions must be manually escaped:

```
<title>
{!SUBSTITUTE(SUBSTITUTE($Request.title,"<","&lt;"),">","&gt;")}
</title>
```

Security Tips: http://wiki.developerforce.com/index.php/Apex_and_Visualforce_Security_Tips

iSEC PARTNERS

Force.com Visualforce

- Tags in the **apex:** namespace escape text contents, eg:

```
<apex:outputText>{!$CurrentPage.parameters.attackedParam}
</apex:outputText>

<apex:page >
<a href="{!$CurrentPage.parameters.evilValue}">vuln link</a>
</apex:page>
```

Security Tips: http://wiki.developerforce.com/index.php/Apex_and_Visualforce_Security_Tips

XSS Lessons

- The PaaS frameworks are status quo compared to classic frameworks.
- Force.com has legacy S-Control, despite fresh start.
- Dev's need to understand language context issues.
  - Or simpler: Don't place parameters in JavaScript, ever.

Context is Key

- There are many unaccounted for contexts:

```
<a onclick="show_content('{{greeting.content|escape}}')">
```

Context is Key

- There are many unaccounted for contexts:

```
<a onclick="show_content('{{greeting.content|escape}}')">
```

JS event handler

Escape for wrong context.

iSEC PARTNERS

Context is Key

- There are many unaccounted for contexts:

```
<a onclick="show_content('{{greeting.content|escape}}')">
```

JS event handler

Escape for wrong context.

- If the greeting content is:

```
‘ + alert(‘xss’) + ‘
```

- The expansion becomes:

```
<a onclick="show_content('&#39; + alert(&#39;XSS&#39;) + &#39;')“>
```

iSEC PARTNERS

SQL Injection

- Requires developer awareness, like XSS.
- New challenge for Cloud vendors… Why?
  - They don't use SQL.
- Departing from the past…
  - Will they reinvent a broken wheel?

GAE GQL Protection

- AppEngine introduces Google Query Language (GQL)
- Examples:

```
query = GqlQuery("SELECT * FROM Song WHERE composer
    = 'Lennon, John'")

query = GqlQuery(
    "SELECT __key__ FROM Song WHERE composer = :1",
    "Lennon, John")

query = GqlQuery(
    "SELECT * FROM Song WHERE composer = :composer",
     composer="Lennon, John")
```

Good code from: http://code.google.com/appengine/docs/python/datastore/gqlqueryclass.html

iSEC PARTNERS

GAE GQL Protection…

- Queries can be composed with explicit APIs:

```
query.filter('title =', 'Imagine').order('-date').ancestor(key)          ‡
```

- Naughty string composition is possible:

```
query = GqlQuery(
    "SELECT * FROM Song WHERE composer = “ + evilStr                     ↓
    )
```

- GQL is query only (no updates).

iSEC PARTNERS

‡ **Good code from:** http://code.google.com/appengine/docs/python/datastore/queryclass.html

↓ **Bad code from:** a twisted imagination.

Azure SQL Protection

- Azure provides parameterized queries via .NET.
- The status quo for SQL Injection mitigation.
  - Use parameterized queries at every call site.
  - See: System.Data.SqlClient

SQL Injection Tutorial: http://msdn.microsoft.com/en-us/library/ms998271.aspx

Azure LINQ potential

- LINQ – First class query support in .NET languages.

```
IEnumerable<string> query = from s in names
                            where s.Length == 5
                            orderby s
                            select s.ToUpper();
```

- Prevent injections *while* making the coder's life easier…

LINQ: http://msdn.microsoft.com/en-us/library/bb308959.aspx

iSEC PARTNERS

Force.com SQL Protection

- Force.com relies on APEX language queries.
- Queries are innate to the APEX grammar:

```
public List<Contact> getMyContacts() {
    return [SELECT Id, Name, Account.Name FROM Contact
                        ORDER BY LastModifiedDate DESC LIMIT 10];
}
```

- One less injection surface is a good thing.

See: http://wiki.developerforce.com/index.php/Force.com_Tutorial:_An_Introduction_to_Visualforce

SQL Lessons

- Different platforms take various approaches:
  - Best: Language integration removes attack surface.
  - Good: Use explicit APIs vs. separate language.
  - OK: Promote Parameterized Queries.
- Still possible to do it wrong in all frameworks.
  - Why? Backwards compatibility.

PaaS Summary

force.com

CSRF Config Steps	2 (nonstd)	2	0
XSS Encoding	Ok	Cumbersome	Legacy: Cumbersome New: Lean
SQL Inj. Mitigation	Parameterized	Parameterized	Language Support

- Make security the default for CSRF.
- Make XSS encoders more context-aware.
- Integrate queries into the application language.

iSEC PARTNERS

Infrastructure as a Service

IaaS Concerns

Linux RNG on IaaS

Infrastructure as a Service

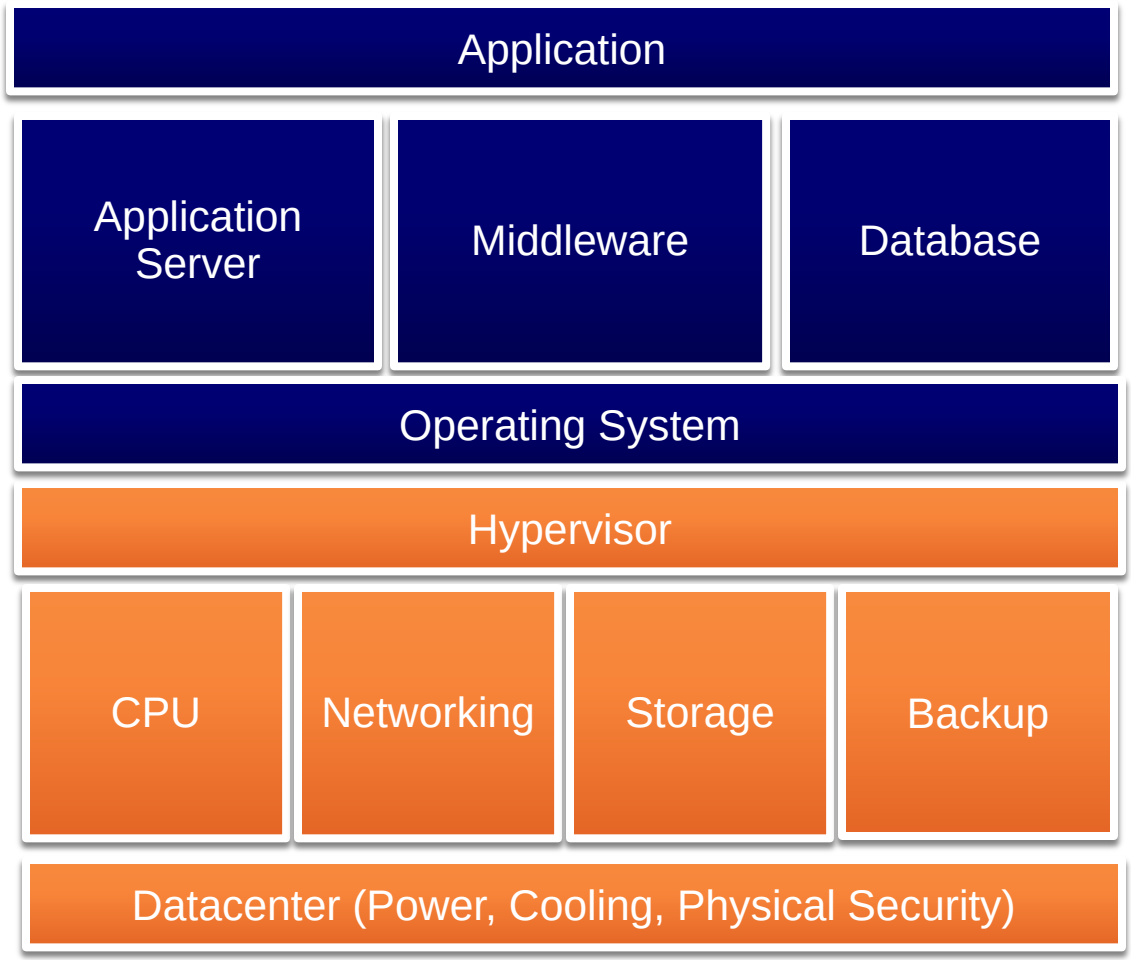

iSEC PARTNERS

IaaS Background

- IaaS is not just virtualization
  - Shorter lived instances
  - Non-persistent local storage
  - Software optimized for cloud lifecycle
  - Often includes helper services like storage

IaaS Concerns

- Flaws in Hypervisor
  - Well researched area, still many bugs to uncover
  - Virtualization bugs are important, but not the last word in IaaS issues
- Services
  - Administrative interfaces can have vulnerabilities
  - Not always accessed over TLS
  - Audit logs are still poor
- Networking
  - "Cheap" IaaS provides = no network segmentation
  - Amazon has ipfilters like ruleset.
  - Generally harder to build secure network

IaaS Concerns – OS Assumptions

- Operating systems aren't built to be cloned at block level

- A lot of unique or secret data
  - Private keys (SSH, SSL, Kerberos)
  - Identifiers (Windows Machine GUID, hostname)
  - Salted password hashes

IaaS Concerns – OS Assumptions

- What else do OSes assume?
  - Running on real hardware
- What does real hardware get you?
  - Performance
  - Maybe DRM benefits
  - TPM
  - Non-deterministic timings

iSEC PARTNERS

IaaS – Random Number Generation

- Quick background:
  - Computers are deterministic, cannot generate random numbers with math.
  - Hardware can extract entropy from physical processes, no standard on x86

- Most random number sources on commodity systems are not **cryptographically** random
  - rand(), rand.nextInt(), random.random()
  - Use predictable algorithms, like the Mersenne Twister
  - Don't use these for online poker

iSEC PARTNERS

IaaS – Random Number Generation

- Modern OSes have PRNGs in the kernel
- Gutterman et. al. "*Analysis of the Linux Random Number Generator*"
- Several issues addressed since this paper, still accurate for entropy gathering

How does Linux gather entropy?

- Input Events
  - From physical keyboard and mouse
  - Do not exist on Xen hosts
- IRQ Events
  - Virtual hardware
  - Not implemented by all device drivers
  - Gutterman dismisses usefulness
- Block device events
  - Estimated 1.03 bits/event on HD
  - Disk numbers predictable
  - Xen implements as ring buffer

iSEC PARTNERS

Linux PRNG Lifecycle

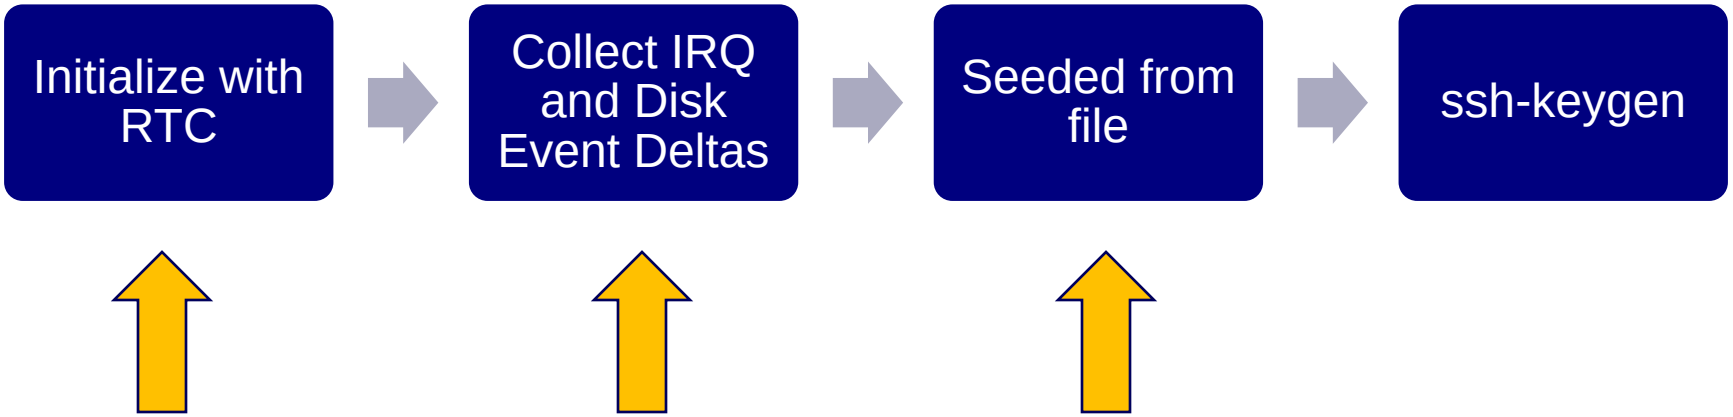

Predictable, shared across the host.

See Goldberg and Wagner

Most IRQs from initialization. Little jitter due to virtualized devices.

random.seed is available on public AMIs

iSEC PARTNERS

IaaS – RNG Vulnerabilities

- Attack lifecycle:
  1. Fingerprint remote EC2 victim to determine likely AMI
  2. Pull down AMI image using S3
  3. Grab AMI's random.seed
  4. Rebundle AMI with instrumented kmod
  5. Run AMI several hundred times to get IRQ/Disk timings
  6. Use PRNG simulator with:
     1. Estimated initial RTC
     2. Most likely IRQ Disk Deltas
  7. Simulate ssh-keygen. Test against fingerprint.
  8. If fail, permutate PRNG. Goto 6.

iSEC PARTNERS

IaaS – RNG Vulnerability

- Is this a practical attack? Maybe.

- Or maybe not:
  - Need estimated "birthdate". Possible over network?
  - Amazon filters SSH ports by default.
  - Would be better if we could optimize testing ssh keys

- Dangerous enough to fix. Mitigations:
  - Xen device driver
  - Custom random.seed
  - Waiting to generate keys

iSEC PARTNERS

Conclusion

1. SaaS takes away many traditional IT controls
2. Incident response on the cloud can be difficult
3. Legal issues are an important stumbling block. Get a good IP lawyer.
4. PaaS vendors can do more to secure web apps on their systems.
5. IaaS leads to unpredictable issues

Bottom Line: State of research into basic technologies does not provide for confident security analysis.

iSEC PARTNERS

Thank you for coming!

Want a copy of the presentation/tool?

Email:

blackhat@isecpartners.com

…and instantly receive all iSEC BH presentations and tools

abecherer@isecpartners.com
alex@isecpartners.com
nathan@isecpartners.com

iSEC PARTNERS