Web and Cloud Based Security

Charlie Britton

September 9, 2024

Contents

- Dev ops Mika Bastrom
- Ethical Hacking bu Philip Blake
- Cyber Recruitment by Chris Philips
- Interview with Edsger W. Dijkstra

The reading list is available and can be followed for the first 5 weeks without going to lectures if we want to. This can be found on the wiki pages.

Chapter 1

The Basics

Vevox Strikes Again

The lecture divulged into rather a lot of Vevox polls. This didn't really achieve much in my opinion but I've left what notes I took as posterity.

- Broadly, we don't want to have to worry about our online privacy and security.
- People are generally happy to spy on enemies but do not want domestic surveillance.
- Corporations spy on us more than governments do because they have many more resources and need our data in order to provide us targeted and relevant advertising.
- APTs were mentioned and their difference to kill chains was highlighted.
- Conventional way of manually designing systems with person not considering security is outdated. Now have to account for the network, code and the people in the system withot necessarily offloading it to the security team.

1.1 Cloud Services

Lecture 1: Cloud Services

2022-10-04T11:00

There has been a paradigm shift over the past few years, with the advent of cloud providers who are able to provide anything from a piece of software to the raw hardware as a service, allowing companies to offload their compute and focus more on getting their job done.

NIST defines cloud computing as the following:

> **Definition 1.** A model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.

Cloud computing provides many benefits to the end users. Key differences between these are highlighted below:

On premises:

- Manually Provisioned
- Dedicated Hardware
- Fixed Capacity
- Pay for Infrastructure Capacity
- Capital (Initial) and Operational Expenses
- Managed by Sysadmins

Cloud infrastructure on the other hand has the following advantages:

- Self-provisioned
- Shared Hardware (Economies of Scale)
- Elastic Capacity
- Pay for Use
- Operational Expenses Only
- Managed via APIs

Cloud providers are typically very competetive with each other, meaning that it is typically cheap enough to use a provider. The five essential characteristics of cloud computing are outlined below.

1.1.1 Essential Cloud Computing Characteristics

On-Demand Self-Service

Because this process is completely automated, the end user is abstracted from the implementation of the new machine, with delivery being orders of magnitude faster than buying, installing hardware and provisioning machines.

Services typically provide a user-friendly UI where end users can spec out their machine to their liking and are up and running within minutes, saving development and sysadmin time.

Shared/Pooled Resources

The resources are typically used from a common pool of compute, with many servers being virtualised and with compute and storage separated so they can grow independently and at massive scale. These infrastructures are designed to be as cost efficient and time efficient to run as possible.

Broad Network Access

Providers will typically have open APIs, able to be controlled using HTTP REST, allowing users to provision and program servers from anywhere that they have an internet connection.

Scaleable and Elastic

As there is a large pool of shared resources that are allocated to the users, providers are able to allow users to create new resources from the pool to use when needed and providers are able to switch machines on and off dependent on global demand instead of leaving the infrastructure all on all of the time.

This process is fully automated and intelligent.

Usage Metering

Service usage is able to be metered according to bandwidth, RAM usage (e.g., in GB/s) or the cost to store something for a certain period of time. Users only pay for the services they use and providers typically have financial incentives for the users to either reserve compute (good for a baseline demand) or to use a different storage service if their data is archival and thus highly unlikely to be accessed.

1.1.2 Deployment Models

Cloud services are categorised into four broad pools:

- **Public** – Infrastructure is available to the general public, owned by the organisation selling cloud services (e.g., AWS).
- **Private** – This is infrastructure for a single organisation, managed either by the organisation or a 3rd party provider and may be hosted on or off-premises. For example, UoS has a data centre off site which we run and host most of our services on. Other backups are in other data centres which we typically co-locate in and just have a few racks of hardware.
- **Community** – This infrastructure is shared by several organisations that have a shared concern and this is either jointly managed or managed by a 3rd party. For example, shared super compute clusters between universities or pharmaceutical companies.
- **Hybrid** – This is just a combination of the above models.

1.1.3 Service Delivery Models

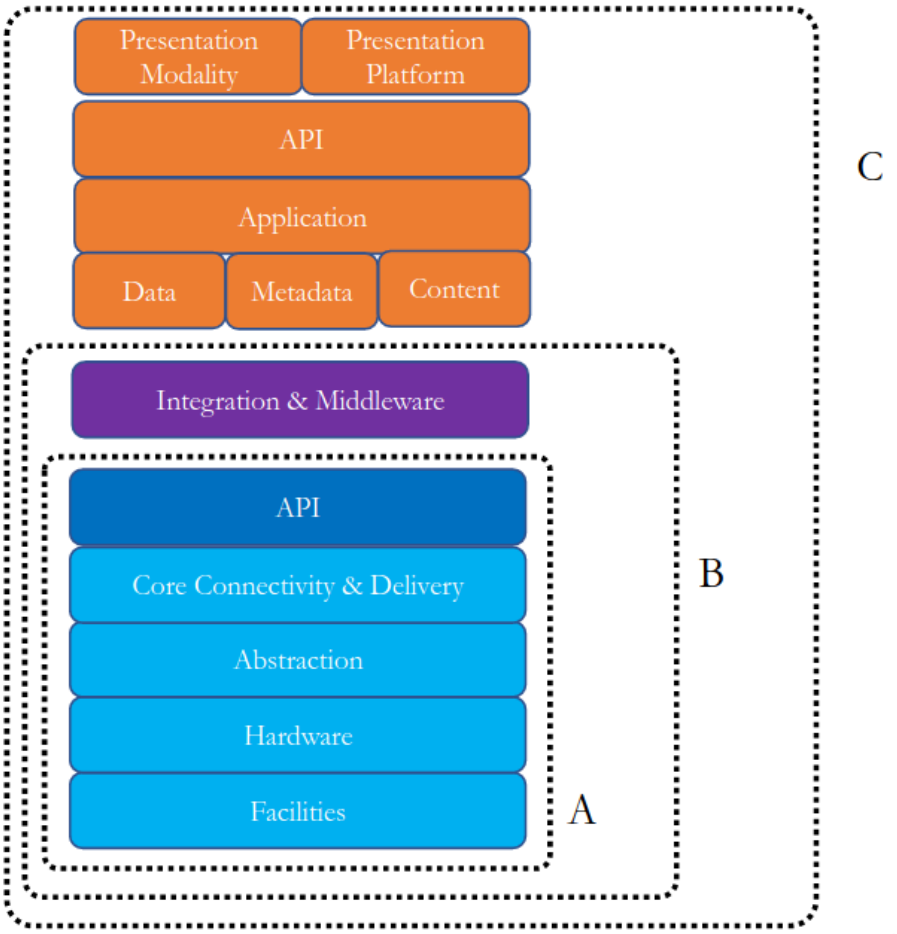

Figure 1.1: CSA Cloud Reference Model. A: IaaS, B: PaaS, C: SaaS

Services can be delivered in different granularities. Just as with the rest of CS, an abstraction can be useful in many cases and the right level of abstraction is a task left to the developer. These different service types are outlined below:

- **IaaS** – Infrastructure as a service. This is the most granular model and the consumer can provision the computing resources within the provider's infrastructure. They can then run arbitrary software and use the VM just as any other computer.
- **PaaS** – Platform as a service. The consumer can create custom applications using the programming tools the provider supports and deploy them onto the provider's services. For example, this may include elastic scaling

according to demand or automatically running CI pipelines.

- **SaaS** – System as a service. This is where the application is built for the cloud and makes use of PaaS.

1.2 Landscape and Policy

Lecture 2: Landscape and Policy

2022-10-10T11:00

If something is measurable then it is a policy.

Restricting the use of calculators is not a policy because it is not measurable but is something that invigilators will have to work around. Something like a person having to mark exam scripts is measurable and so can be a policy.

Definition 2. A policy is a statement that reflects the rules that everyone should abide by. All good policies are organised as follows:

- Purpose
- Policy Compliance
- Date Tested
- Date Updated
- Contact

1.2.1 Models and Frameworks

A model is a step by step guide to follow, a framework provides a foundation that we can use to build the model but we don't typically use all the pieces of a framework.

Security Frameworks

Several frameworks exist for information security online. We have ISO 27001, NIST and BSI. These are important because to work with certain entities (e.g., governments), we need to meet minimum compliance standards.

The government came out with the cyber essentials scheme, which is easy for all companies to understand and it is easy to be compliant with the government. SMEs may be able to be compliant with the cyber security essentials and a larger corporation can abide by ISO 27001.

Security Models

Models include IoT, blockchain and clouds. We know with IoT that we have sensors, edge nodes, collectors and processing computers. Blockchain has different types of nodes that need to be put in the infrastructure.

To measure the performance of an IT system, we need goals, questions and metrics. A framework only puts everyone at the same level. Some people need to be more secure than others because their computers are more of a target.

1.2.2 Security Scales

There are 2 scales for measuring the security of the system.We have the SANS Endpoint Security Maturity Model Curve, which has 5 different scales of how an organisation protects their assets:

1. Random or disorganised
2. Reactive or tactical
3. Preventive
4. Organised or directed
5. Proactive, comprehensive, continuous and measurable

The second model, the Cybersecurity Capability Maturity Model (C2M2) Maturity Indicator Levels (MILs) are more catered to devops. These are as follows:

- **MIL0** – Not accomplishing objectives, or accomplishing with manual process
- **MIL1** – Accomplishing objectives, but with some automation, but minimal or ad hoc process
- **MIL2** – Established and followed standard operating procedures, more automation
- **MIL3** – Mature implementation with high degree of automation and highly optimised

1.2.3 Building Reference Diagrams

First, we define hosts, then networks and finally stakeholders. Hosts are basically any computer that is networked and is not part of an industrial control system. Network assets are any asset that enables communication, for example routers and switches. Stakeholders are persons involved in the day to day operation of the organisation.

Reference diagrams are not typically given out from a company or exist within the company. We therefore work with the client to ensure that the network map is as comprehensive as possible and that we are covering all of the hosts and the whole network.

We can illustrate the diagram with possible threats, showing what assets, stakeholders and network vulnerabilities there are. If we want to go more in depth,

we can show 'forbidden relationships', which are things that shouldn't be allowed to interact (e.g., sales staff having access to the original schematics or diagrams for the system).

We can then add monitoring nodes and identify host assets that may come under significant load. It is also good to identify persistent assets.

1.3 Infrastructure as Code

Lecture 3: Infrastructure as Code

2022-10-11T11:00

Devops engineers are a combination of both a developer and network manager. Pure network administrators need to know how to operate switches (e.g., Cisco CCNA). This type of job is now within data centres and so only people in the datacentre really need to understand how to manage switches.

Modern devops engineers need to know about the cloud and CI/CD and much less about the actual hardware they work on. Historically, we needed to deploy our own servers and manage them but now we just program the deployments we need so that we don't have to manually configure each server.

Security testing is a very worthwhile career path to choose as it automates the process of security testing somewhat and is a skill that is needed but underserved in the job market.

1.3.1 Infrastructure Creation Comparison

Automated vs manual provision are the two ways we can deploy infrastructure today. Automatically provisioned infrastructure is:

- Designed using code
- Written using high level programming language
- Can be tested and reviewed like code
- Can be relocated
- Reduces bugs
- Increases development speed
- Allows one employee to manage a big infrastructure
- Allows exact copies to be deployed
- Can be changed easily

Terms

- **Artifact** – what gets versioned, tested and deployed

- **Procedures** – also known as imperative, e.g., a Bash script
- **Functions** – also known as declarative, where desired state is defined, tool conforms the system to the model of desired state.

1.4 Virtualisation

Lecture 5: Virtualisation

2022-10-18T11:00

The use of an FPGA allows businesses to start offering hardware functions such as radar systems, etc. which allow for faster processing. FPGAs allow us to create the hardware acceleration on the fly and then adapt it later.

Using an FPGA to play games and creating virtual machines are both virtualised hardware.

1.4.1 Virtualised Software

Virtualising means to take something that was not meant to be virtual and making it virtual. Google docs was designed to be a SaaS so it is not a virtual software, installing an OS on a virtual machine is a virtualised software.

1.4.2 Virtualisation Types

There are different ways we can virtualise computing resources. Hardware, software, OS, server and storage are all different types.

1.4.3 Virtualisation Tools

There are different tools for different computer types. We can virtualise on the desktop with VMWare workstation or VirtualBox, containerise (which is lighter than full virtualisation) such as Docker or LXC, hypervise on-premises or colocated servers with vCenter Server Appliance and ESXi or manage containers with containerd, K8s or VMWare integrated containers.

Hypervisors help us to virtualise bare metal machines. They can however have vulnerabilities which allow us to escape the virtualised machine and then control the machine that runs the VM.

Hypervisors can also be hijacked and a malicious fake hypervisor can be installed and manage the entire server system.

Meltdown is the other main vulnerability in a hypervisor, with the fundamental separation of user applications being separated being broken and allowing access to the RAM of another VM.

Specter breaks the isolation between different applications. Attackers trick programs following best practices into leaking their secrets. Following the safety

checks of best practice increases the attack surface and make applications more susceptible to a Specter attack.

Lecture 6: Docker

2022-10-24T11:00

When developers are developing an application, they need to share a common environment. Setting up a VM potentially takes a long time to do and is quite heavyweight.

Docker allows us to have a consistent environment shared across machines.

Docker containers have a mini kernel which is the basics that the container needs to operate but not a full OS. The docker daemon runs on the host OS and the docker container interfaces with the host.

Containers include config files, processes and dependencies and networking but doesn't have a full operating system.

Docker has a large potential vulnerability where the container is escaped and we gain control of the host system.

Coursework no active scanning just create a list and a justification for what we have picked. Choose what seems to be reasonable. Most of the job of a security analyst is to guess what they have.

1.5 Service Level Agreements

Lecture 7: Service Level Agreements

2022-10-25T11:00

Service level agreements are parts of contracts that define what services a provider will provide and the required level or standard for the uptime of those services.

This is usually a part of an outsourcing effort by a company and the agreement seeks to ensure that there are consequences for companies losing their outsources services.

The objectives of a service level agreement are to set out the objectives for the services to be provided. These services will improve performance, save costs, provide access to skills and technologies. These are services that cannot be provided internally.

1.5.1 Acceptable Level of Service

Customers often want the highest level of service possible. This is not feasible for many cloud providers who wish to be competitive and therefore share hardware or have less redundancy.

A provider who fails to provide the service for the agreed SLA will typically include a service credit regime. This compensates for a failure and should be roughly equal to the loss the company suffers.

A good SLA would include a clause allowing the company to terminate the contract if the provider fails to provide the service at the agreed level. The customer should also have the right to sue for damage and the right for service credit.

Lecture 8: Common Cloud Security Threats and Assessment Q&A

2022-10-31T11:00

10 common threats to the cloud

Common threats are explained usigng the cloud reference model, which is similar to the CSA model we have seen earlier.

Accountability and Data Ownership

Consider data toxicity, who owns the data, where to store, how to destroy, how to encrypt.

IDentity providers SAML OpenID Access control design

Regulatory Compliance Data perceived as secure in one region is not assumed secure by another. Lack of transparency in the underlying implementations makes it difficult to demonstrate compliance to owners. There is a lack of consistent clobal standards for handling data.

1.6 Service and Data Integration

1.6.1 Encrypting Data

Lecture 10: Cloud Security Vulnerabilities Part 2

2022-11-01T11:00

When data is in transit, it needs to be encrypted to ensure that only authorised users have access to it. Data is also authenticated to ensure that it is not modified through something such as a MITM attack.

At rest, to maintain the CIA triad, we ensure that the data is encrypted to maintain confidentiality, periodically use checksums and parity data to maintain data integrity and also use hot spares and a clustered filesystem if we want to maintain availability.

1.6.2 Multi Tenancy and Physical Security

This is where multiple companies or departments have data hosted on shared infrastructure. We need to be careful to ensure that we maintain adequate logi-

cal separations, keep tenant data separate from each other and avoid malicious or ignorant tenants (who may compromise data).

Furthermore, the more tenants that use a shared service, the more a single point of failure becomes an issue. This can fail if there is uncoordinated changes to configurations, bringing multiple tenants offline and causing much more widespread issues.

Furthermore, with more tenants, we get much higher performance risks than with each separated from the others. Attacks on this infrastructure include cross-tenant attacks, side channel attacks, scanning other tenants and denial of service.

To overcome these risks, design for multi-tenancy needs to be built in from the start, with encryption per-tenant, controlled and coordinated change management, an ability to audit the administrative access of the provider and a framework for third party assessment of the services.

1.6.3 Incident Analysis and Forensic Support

Following a data breach, we need to be able to identify and manage critical vulnerabilities to respond to the incident quickly and effectively. Cloud computing can make responding to these incidents more difficult because audit data and events may be logged across many data centres.

When designing for the cloud, we want to look at their policy on handling, evaluating and correlating event logs across data centres. We also should make sure that they have technologies such as machine imaging in place so that we can forensically analyse security incidents.

1.6.4 Non Production Environments

When developing concepts and tests, we should use locally hosted and not widely available locations. We should use MFA and treat any pre production environment and data as top secret.

1.7 Web Security

Lecture 11: The Security of 'the Web' & Vulnerabilities: The Big Three

2022-11-14T11:00

OWASP tracks a list of the top vulnerabilities.

1.7.1 XSS

We can inject new HTML into pages in a way that can be executed by other users. For example, crafting a link or set of steps so that we get the server

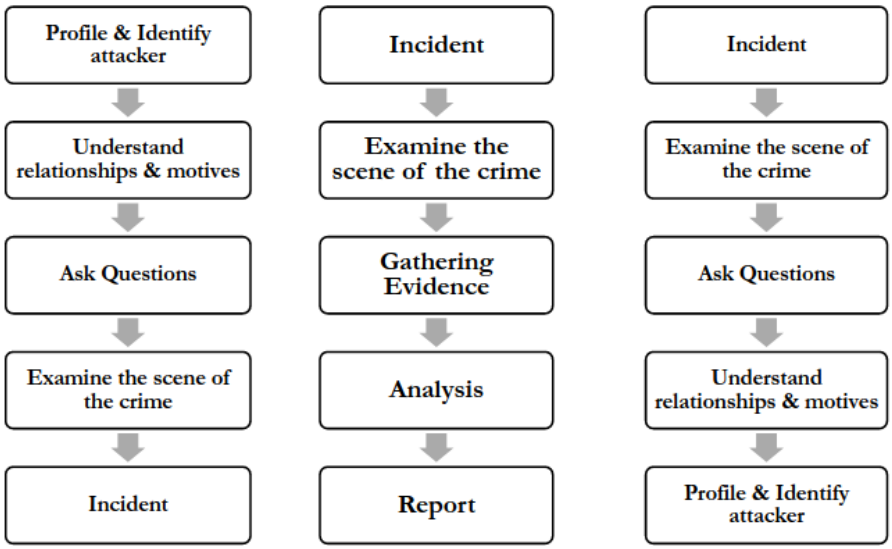

Figure 1.2: Investigator Teams. Left to right: Incident Management, Physical Forensics, Digital Forensics

to give us code we inject. For example, something that returns the raw text provided (which could be HTML/JS)

Lecture 12: Web Security

2022-11-21T11:00

1.7.2 Scoping

We agree on what is being tested, when and how. This is to define the boundaries and ensure that we are doing the job properly and in a way that the company will increase.

We can then gather information from one of their preproduction systems and start to identify vulnerabilities. We then try and exploit the application to eliminate false positives. We then do the post-exploit steps (What can we do now), then the report.

When doing the coursework, we want to be systematic and logical about what we are attacking, and plan it instead of trying lots of different URLs at random. We want to be careful with the rate we attack the server as otherwise we can trigger the firewall.

1.7.3 Burp Suite

This is an intercepting proxy. It collects all of the data that goes over the network to allow us to track what we have done. We can then review the requests and responses the server gives us and build a map of the network.

This used to be simple, but with the advent of TLS, we have to add local certificates to proxy traffic properly.

Set limits of 2 concurrent, 300ms delay with random variations.

This wants to be done on a VM, where browser settings are changed, then browser isn't secure.

1.7.4 SQLMap

This is an application that allows us to try and use a database as if we were logged in through an exploit. For example, in Rob the Shop, we could SQLMap the user or password parameter. This should be done after a couple of manual executions.

1.7.5 dirbuster

This allows us to concatenate words from a wordlist into directory trees and see if we get a response.

Want to create wordlists that are short and specific to the application we are using, also exclude things we're not interested in.

> Don't want to use Dirb or content discovery with rob with giant lists

Keeping track of where we can enter data is something rolled into burp suite