IoT Security

ΕΠΛ 428: IOT PROGRAMMING

Dr. Panayiotis Kolios
Assistant Professor, Dept. Computer Science,
KIOS CoE for Intelligent Systems and Networks
Office: FST 01, 116
Telephone: +357 22893450 / 22892695
Web: https://www.kios.ucy.ac.cy/pkolios/

Πανεπιστήμιο Κύπρου

- From smart grids, to intelligent transportation systems, security of the networks, devices, and the applications that use them is critical
- Among the very few disciplines that works against desired outcomes
- To further complicate matters, these external forces are able to leverage traditional technology as well as nontechnical methods (for example, physical security, operational processes, and so on) to meet their goals.
- Many potential **attack vectors**
    - can result in physical damage and impact human lives environment, and infrastructure.

- Attacking is easy
  - Criminal organizations and nation-states
    - May have plenty of resources, and powerful capabilities
  - Attackers can ignore ethical concerns
  - Attackers need to find ONE way in
- Defending is hard Defenders must protect on ALL fronts
  - Defences interfere with business goals
  - Hard to enforce laws beyond national barriers
    - Especially when attacks come from states unable or unwilling to cooperate
  - Targets: devices running vulnerable software
    - Hard/impossible to find all the bugs in a piece of code
      - Not just bugs, but also design and logical flaws
    - Software license agreements do not hold vendors accountable

- Delivery with USB stick (no internet connection necessary)
  - Worm that spread on Windows machines aiming to reach specific Siemens controllers (ICS, and PLCs connected to variable-frequency drives Exploited 4 zero-day flaws)
- Replay measurements to control center and execute harmful controls
  - Would operate controller in a way to damage uranium-enriching centrifuges
- First example of cyber weapon creating serious physical damage

"The Real Story of Stuxnet", IEEE Spectrum, 2013

The New York Times

Science

WORLD U.S. N.Y. / REGION BUSINESS TECHNOLOGY SCIENCE HEALTH SPORTS OPIN
ENVIRONMENT SPACE & COSMOS

Malware Aimed at Iran Hit Five Sites, Report Says

By JOHN MARKOFF
Published: February 11, 2011

The Stuxnet software worm repeatedly sought to infect five industrial facilities in Iran over a 10-month period, a new report says, in what could be a clue into how it might have infected the Iranian uranium enrichment complex at Natanz.

TWITTER
LINKEDIN
PRINT
REPRINTS

Πανεπιστήμιο Κύπρου

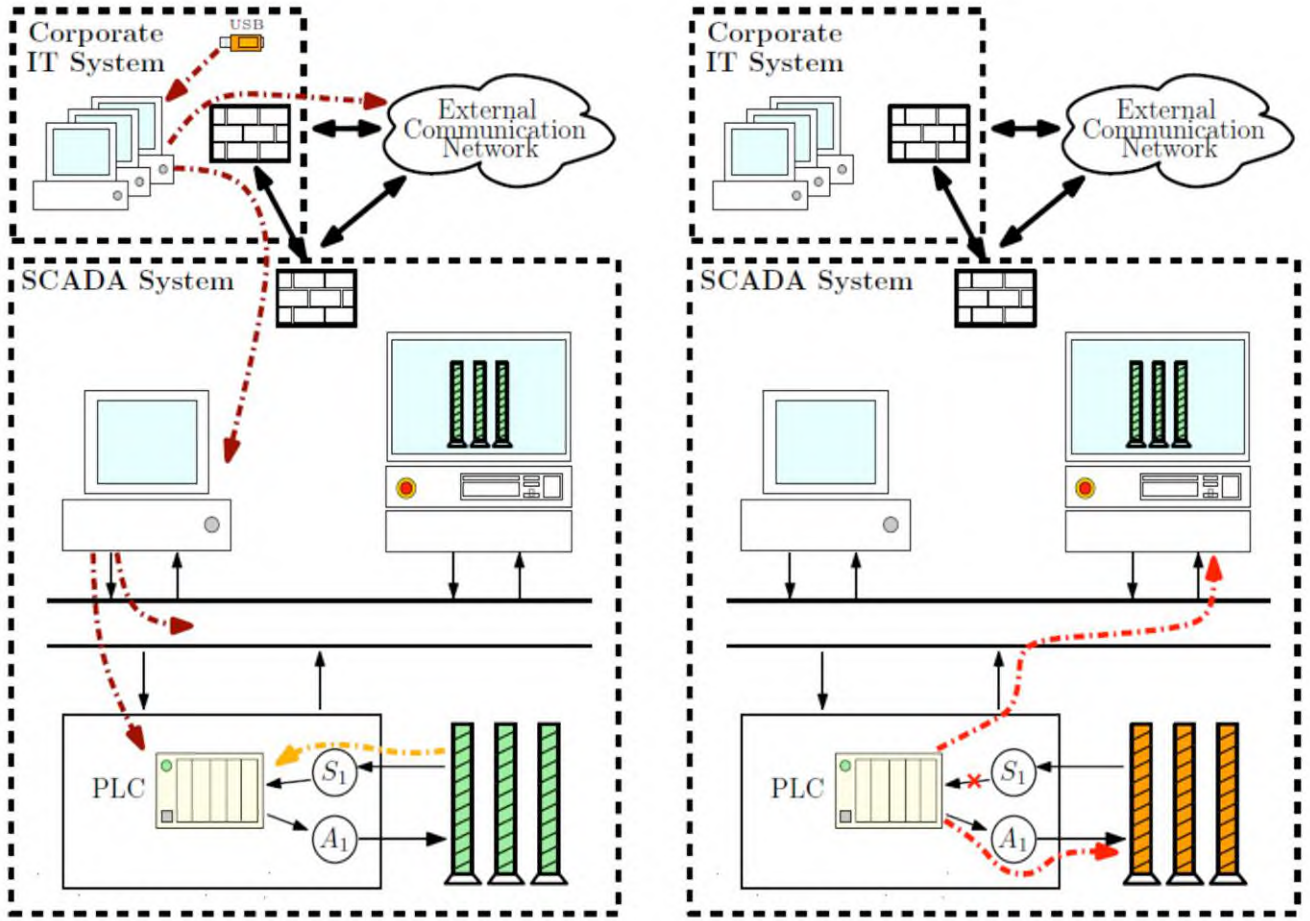

(a) Infection and data recording.

(b) Covert sabotage.

Πανεπιστήμιο Κύπρου

- BlackEnergy(2007-)
- From arstechnica.com:
  - “In 2014 ... targeted the North Atlantic Treaty Organization, Ukrainian and Polish government agencies, and a variety of sensitive European industries”
  - “booby-trapped macro functions embedded in Microsoft Office documents”
  - “render infected computers unbootable”
  - “KillDisk, which destroys critical parts of a computer hard drive”
  - “backdooredsecure shell (SSH) utility that gives attackers permanent access to infected computers”
  - More advanced, more autonomous, follow-up attack in 2016: “Crash Override”

- In 2000, the sewage control system of Maroochy Shire in Queensland, Australia, was accessed remotely, and it released **800,000 liters of sewage** into the river and coastal waters
- Actor has been a consultant on a water project, conducted the attacks after he was refused a full-time job with the Maroochy Shire government
    - **accessed computers controlling** sewerage system
    - altered **electronic data of pumping stations** and causing malfunctions in their operations
    - 150 stations pumping sewerage to treatment plants
    - each pumping station had installed an RTU computer capable of receiving instructions from a central control centre, transmitting alarm signals and other data to the central computer and providing messages to stop and start the pumps at the pumping station
    - communications between pumping stations and between a pumping station and the central computer were by means of a private **two-way radio system**

- Cyber incidents in critical infrastructures in the US (Voluntarily reported to ICS-CERT)

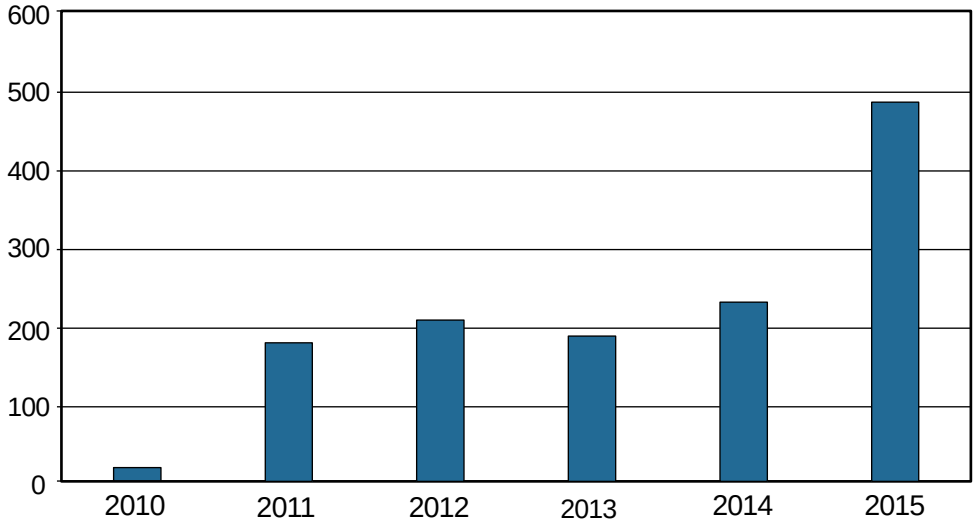

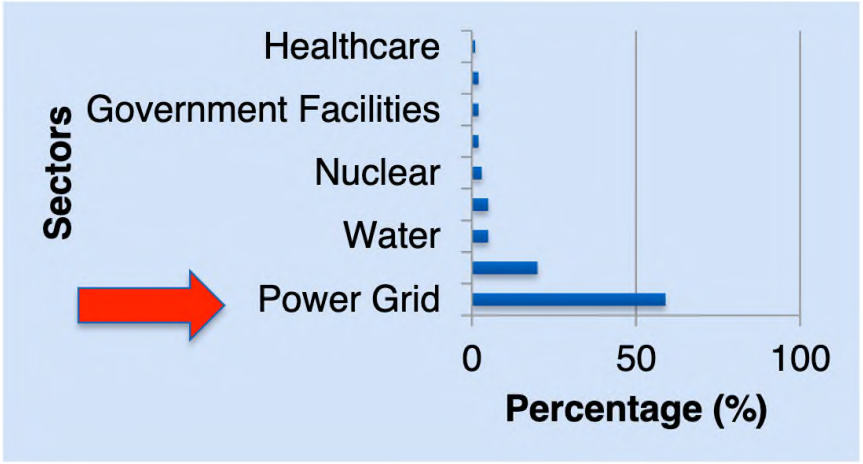

https://ics-cert.us-cert.gov https://www.risidata.com/

Πανεπιστήμιο Κύπρου

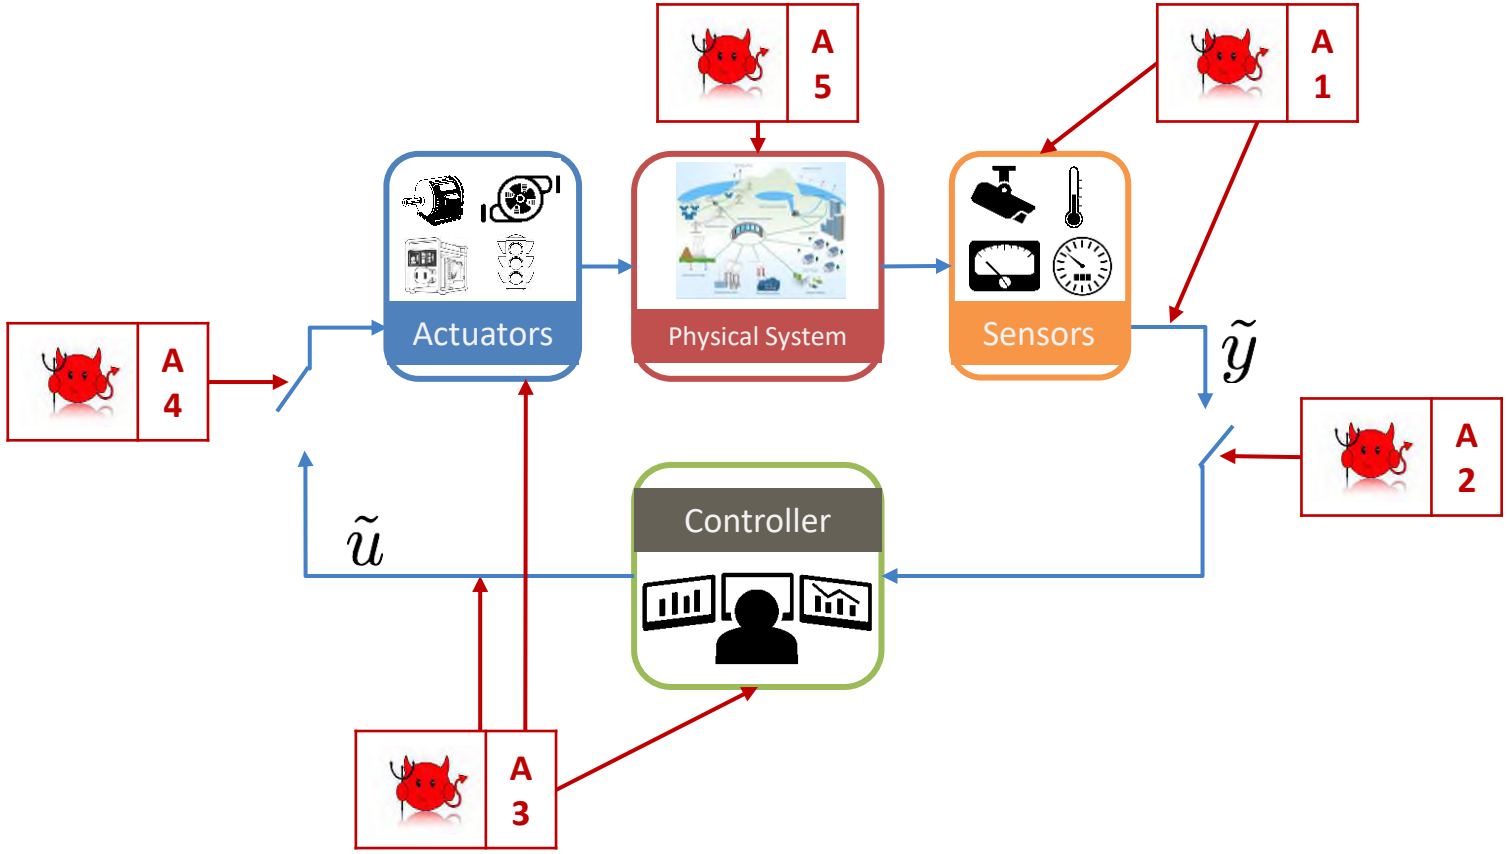

- Deception attacks (A1, A3, A5) may affect the monitoring and control of physical processes,
- Denial of service attacks (A2, A4) may disrupt the proper transfer of information between the components

- Vulnerabilities are software bugs that attackers can exploit in order to compromise computers
- Exploits are pieces of software that take advantage of a vulnerability in order to access or infect a computer
- A zero day vulnerability/exploit is one that is unknown to the software vendor
- Who finds zero-days, what do they do with them, and why?
  - Vendor's employees: fix (it's their job)
  - Security companies: sell, disclose (it's part of their business model)
  - Independent security researchers: sell, disclose (for profit or fame)
  - Academics: disclose (to make the world a better place)
  - Government agents: exploit, disclose ("to protect and to serve")
  - Criminals: exploit, sell (for profit)
  - Terrorists: exploit (to wreck havoc)
  - Hacktivists: exploit (to make the world a better place, according to them)

- **Spoofing**: pretending to be something/somebody else
- **Tampering**: modifying without permission
- **Repudiation**: denying to have done something
- **Information Disclosure**: revealing information without permission
- **Denial of Service**: prevent a system from providing a (timely) service
- **Elevation of Privilege**: achieve to do more than what is intended

- Some threats may belong to more than one category

- Different approaches to evaluating threats
- Beware of formulae that quantify risk
  - It's difficult to estimate realistic parameters
  - Companies don't release breach data, although this is changing
  - Black Swan problem: extremely rare events are hard to predict and quantify
- **DREAD**
  - Score each threat between 5 (lowest) and 15 (highest)
  - Designed at Microsoft, now used in other companies

	Rating	High (3)	Medium (2)	Low (1)
D	Damage potential	The attacker can subvert the security system; get full trust authorization; run as administrator; upload content.	Leaking sensitive information	Leaking trivial information
R	Reproducibility	The attack can be reproduced every time and does not require a timing window.	The attack can be reproduced, but only with a timing window and a particular race situation.	The attack is very difficult to reproduce, even with knowledge of the security hole.
E	Exploitability	A novice programmer could make the attack in a short time.	A skilled programmer could make the attack, then repeat the steps.	The attack requires an extremely skilled person and in-depth knowledge every time to exploit.
A	Affected users	All users, default configuration, key customers	Some users, non-default configuration	Very small percentage of users, obscure feature; affects anonymous users
D	Discoverability	Published information explains the attack. The vulnerability is found in the most commonly used feature and is very noticeable.	The vulnerability is in a seldom-used part of the product, and only a few users should come across it. It would take some thinking to see malicious use.	The bug is obscure, and it is unlikely that users will work out damage potential.

- Recommend a response: META
- **Mitigate**: make a threat harder to exploit
  - Threat: spoofing via password brute-forcing
  - Mitigations:
    - Require longer, more random passwords
    - Lock account after 3 failed attempts
    - Use biometrics instead (too expensive?)
- **Eliminate**: typically, remove the feature that was exposed to the threat
  - Longer passwords don't eliminate spoofing
  - Giving up on user accounts does (clash with business objectives?)
- **Transfer**: let another party assume the risk
  - We still want user accounts: "Log in with Facebook"
  - Cost: Facebook gets info about your customers
- **Accept**: when other options are impossible or impractical
  - Nothing can prevent a lucky hacker from guessing a password on first try
  - Important to keep track that the threat remains valid
- Cost-benefit analysis of each response depends on business objectives
- Document your response: a good way is to use a bug reporting system

- In IoT convergence between OT and IT
- IoT processes optimized for specific operation
- Encompass security and networking requirements for a control system using a logical framework
  - **Purdue Model for Control Hierarchy**
  - Segments devices and equipment by hierarchical function levels and areas
  - Has been incorporated into the ISA99/IEC 62443 security standard

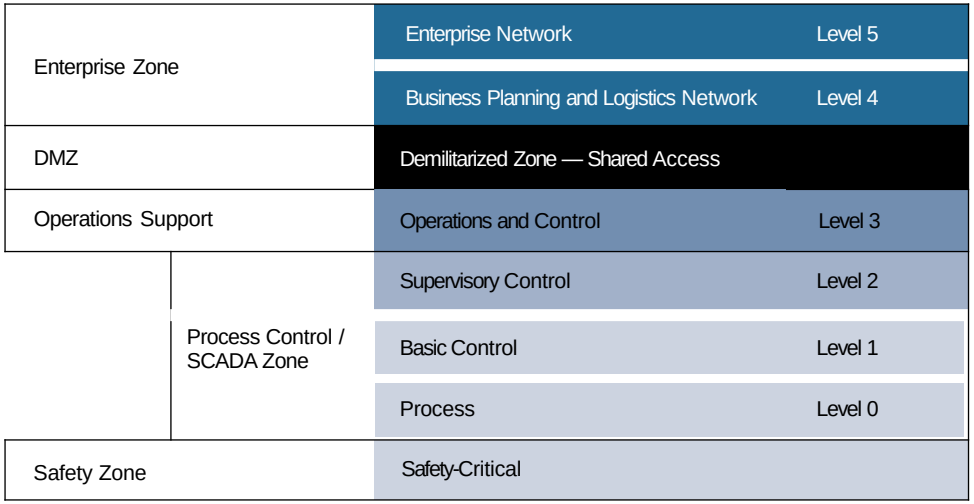

Πανεπιστήμιο Κύπρου

- Enterprise zone
  - **Level 5**: Enterprise network applications such as Enterprise Resource Planning (ERP), Customer Relationship Management (CRM), document management, and services such as Internet access and VPN entry from the outside world
  - **Level 4**: Business planning and logistics; IT services including scheduling, planning systems, and security monitoring
- Industrial demilitarized zone
  - **DMZ**: A buffer zone where services and data can be shared between the operational and enterprise zones. Distribute organizational control. No traffic should traverse the DMZ
- Operational zone
  - **Level 3**: Operations and control: Functions involved in managing operations, monitoring and controlling the entire system; system control, security management, network management
  - **Level 2**: Supervisory control: Zone control rooms, controller status, control system network/application administration, and human-machine interface (HMI) and historian
  - **Level 1**: RTUs, Controllers dedicated HMIs, and other applications may talk to each other to run part or all of control functions
  - **Level 0**: Process devices such as sensors / actuators and machinery
- Safety zone
  - **Safety-critical**: This level includes devices, sensors, and other equipment used to manage the safety functions of the control system

- Seem largely similar to those in traditional IT environments
- BUT with more profound impact
  - IoT devices reveal more sensitive information to the system for processing
  - Information may now include more coverage and depth of personal data
  - IoT may now be able to control physical world
  - Much more complicated to deal with due to involvement of diverse machines & users

- IoT devices are constrained
  - lack the required processing, memory, storage, and power requirements
- Can not support state-of-the-art authentication protocols
  - That are computationally and power hungry
  - Require manual parameter tuning
- Consist of multiple technologies
  - Radio-frequency identification (RFID), virtualization, etc.
  - Each of these technologies has its own vulnerabilities.
  - Must secure the chain of all of those technologies

- Scalability
  - **17 billion** active IoT devices currently
  - Need decentralized security strategies
- Availability
  - Perpetual operation - cost of downtime can far exceed the cost of remediating incidents
- Time-sensitivity

- Confidentiality
  - ensures that the exchanged messages can be understood only by the intended entities.
- Integrity
  - ensures that the exchanged messages were not altered/tampered by a third party.
- Authentication
  - ensures that the entities involved in any operation are who they claim to be. A masquerade attack or an impersonation attack usually targets this requirement where an entity claims to be another identity.
- Availability
  - ensures that the service is not interrupted. Denial of service attacks target this requirement as they cause service disruption.

- Authorization
  - ensures that entities have the required control permissions to perform the operation they request to perform.
- Freshness
  - ensures that the data is fresh. Replay attacks target this requirement where an old message is replayed in order to return an entity into an old state.
- Non-repudiation
  - ensures that an entity can't deny an action that it has performed.
- Forward Secrecy
  - ensures that when an object leaves the network, it will not understand the communications that are exchanged after its departure.
- Backward Secrecy
  - ensures that any new object that joins the network will not be able to understand the communications that were exchanged prior to joining the network.

- Asset Valuation: Identify and assign values to all assets (hardware, software and data) in terms of the impacts of confidentiality, integrity and availability.
- Threat Assessment: Identify and assess the potential threats that can lead to security failure.
- Vulnerability Assessment: Identify and assess the weaknesses in an asset that may be exploited by the identified threats.
- Risk Assessment: Analyse the levels of risks based on the identified threats and vulnerability and the values of assets.

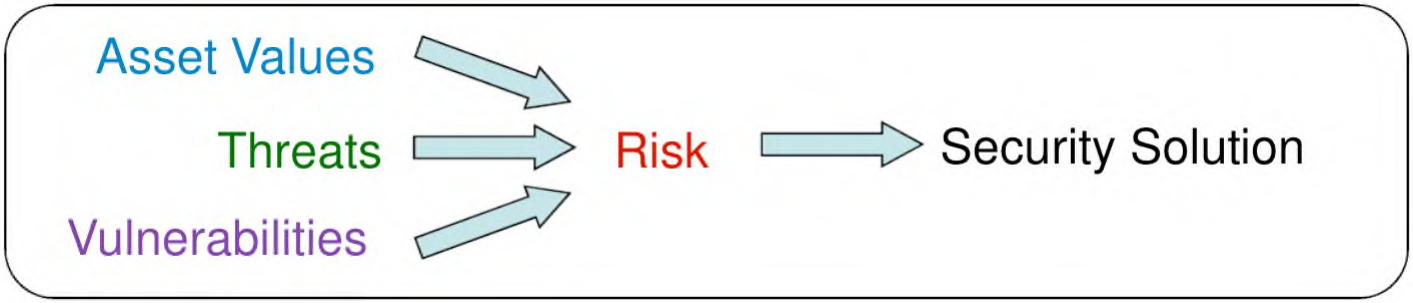

- Information that can be sensitive in IoT:
  - Data generated by sensors (may contain information about physical setup or user behaviours)
- Commands produced by servers (may contain instructions to alter physical environment)
- Information related to IoT services (may contain instructions to control the IoT system)
- Traces stored by the system (may contain meaningful information about users)
- Risks:
  - Confidentiality: unauthorised collection of sensitive information causing privacy leaks
  - Integrity: unauthorised modification of sensor data or actuator commands
  - Availability: blocking data access or services

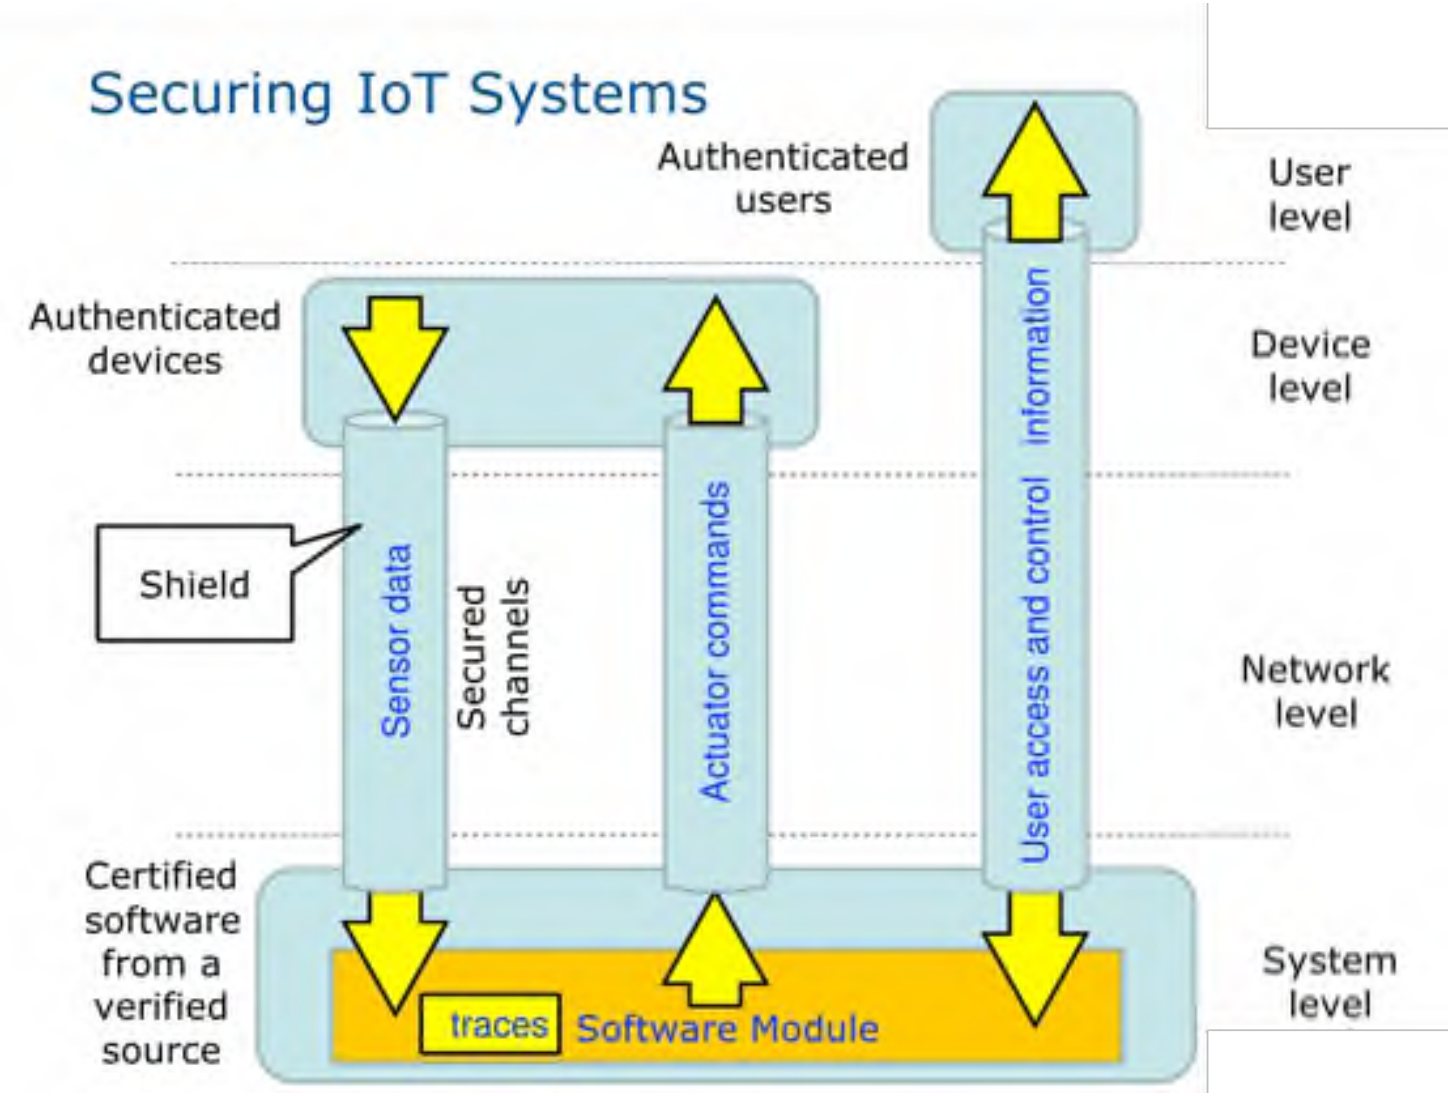

Πανεπιστήμιο Κύπρου

- User level
  - Gain unauthorised access to the system by compromising user access, such as users accidentally reveal their access code by responding to malicious emails, trojian horse, spyware, etc.
- Device level
  - Unauthorised inclusion of devices
  - Tampering of legitimate devices

- Network level
  - Sniffing of data
  - Insertion of data
  - Modification of data
  - Removing of data
  - Interception of traffic flow
  - Jamming channel (interfering the system operation)
- System level
  - Unauthorised access to the hosting servers by exploiting vulnerabilities in the OS or physical machines
  - Access to the software modules and traces by exploiting bugs, flaws in update procedure, misconfiguration, etc.

- Device level
  - Embedded device: tiny and often hidden, remotely deployed, may be tampered without noticing
- Network level
  - Wireless channel: easy to join, sniff data, insert data, jam channel, etc
- System level
  - Server OS: may be compromised
  - Software module: may contain bugs for exploitation
  - System configuration: may be misconfigured
- User level
  - User practice: user access may be compromised

- To prevent unauthorised adding of devices into the network
- Employing authentication procedure to authenticate devices to prevent tampering of devices
- Hard to prevent, but we can monitor possible misbehaviours of devices and avoid storing any sensitive information in the devices
- Security Solution
    - Silicon-based (strongest)
    - Firmware-based
    - Software-based (weakest)

- To prevent sniffing of data
  - Employing strong encryption scheme
- To prevent unauthorised manipulation (insert, modify, delete) of data
  - Employing authentication and enforce communication only with authenticated and well behaved devices
- To prevent interception of traffic flow
  - Employing authentication and enforce communication only with authenticated and well behaved devices
- To prevent channel from jamming
  - Hard to prevent, but we can monitor the channel activities to detect abnormality and introduce backup channels

- Each internet host needs an IP address: 155.198.140.14
- Network services are multiplexed through the same IP address using ports
  - 155.198.140.14:80
  - Common services tend to be hosted on standard ports
    - SSH: 22, DNS: 53, HTTP: 80, HTTPS: 443
- One machine can have multiple IPs
  - Over time: connect at home, at work, on the go
  - At the same time
    - Client with wireless and Ethernet connections
    - Dual-homed host (firewall, gateway)
- Multiple machines may share the same IP
  - Home router connecting desktop, laptop, iPhone
  - Port- or name-based virtual hosting of websites
- Datagram (or packet): headers + payload that is sent as a single unit on the network

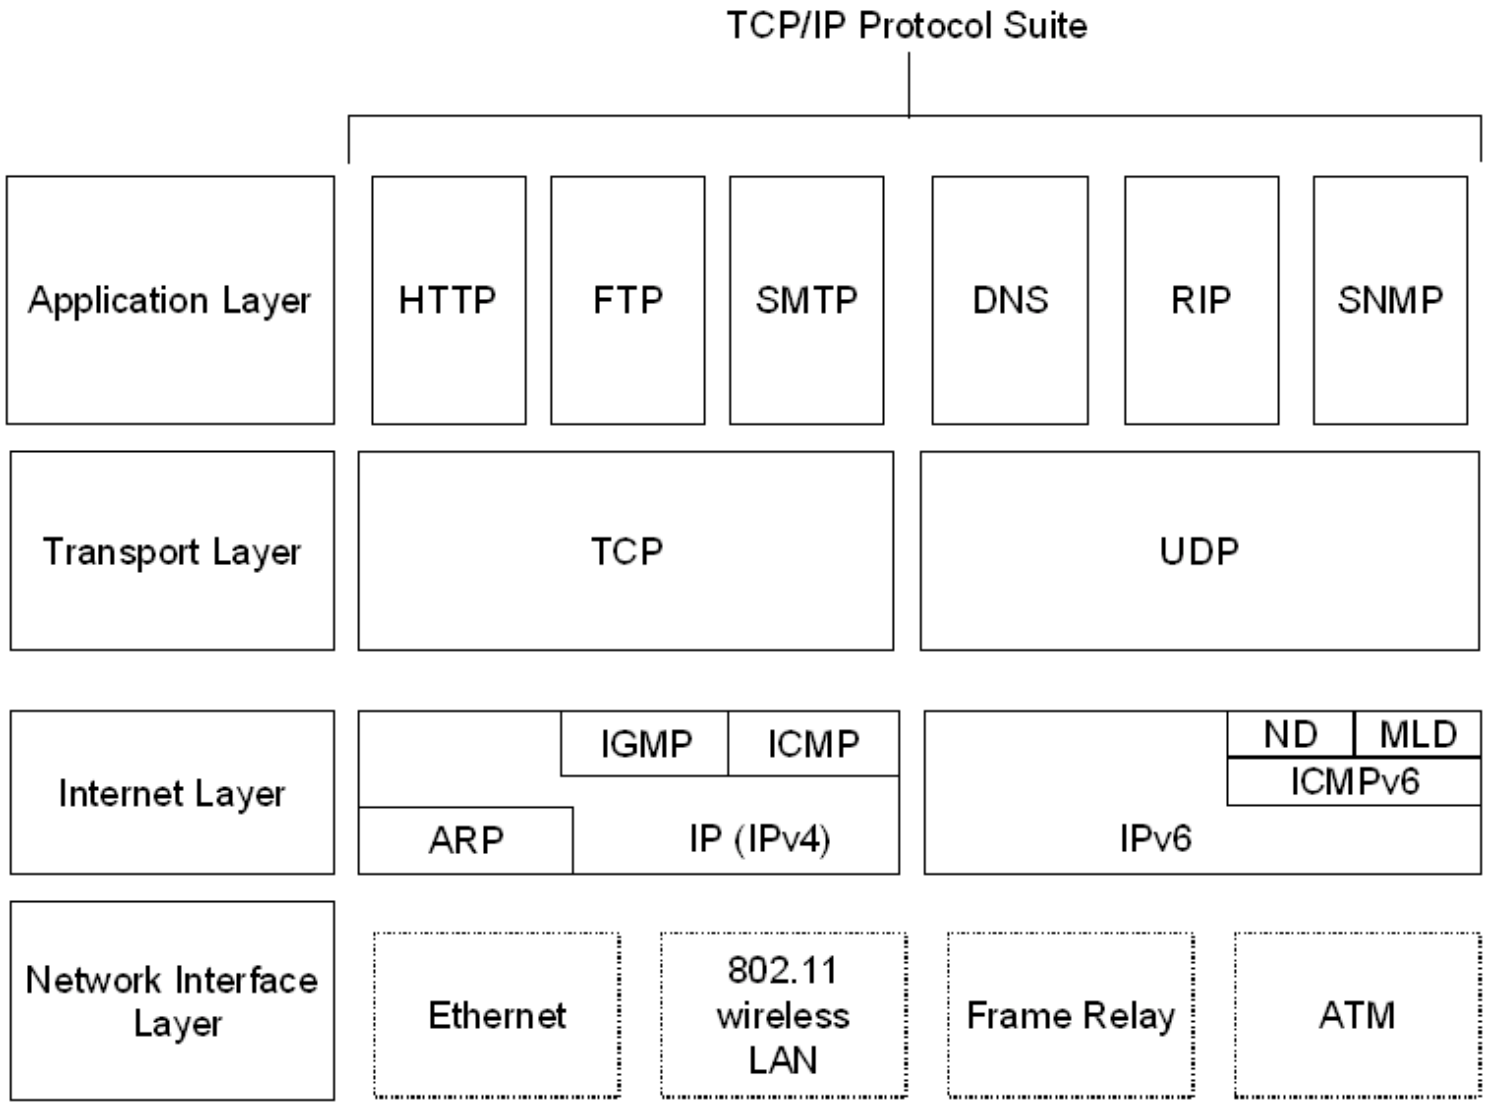
TCP/IP Protocol Suite
Application Layer
HTTP
FTP
SMTP
DNS
RIP
SNMP
Transport Layer
TCP
UDP
Internet Layer
IGMP
ICMP
ARP
IP (IPv4)
ND
MLD
ICMPv6
IPv6
Network Interface Layer
Ethernet
802.11 wireless LAN
Frame Relay
ATM

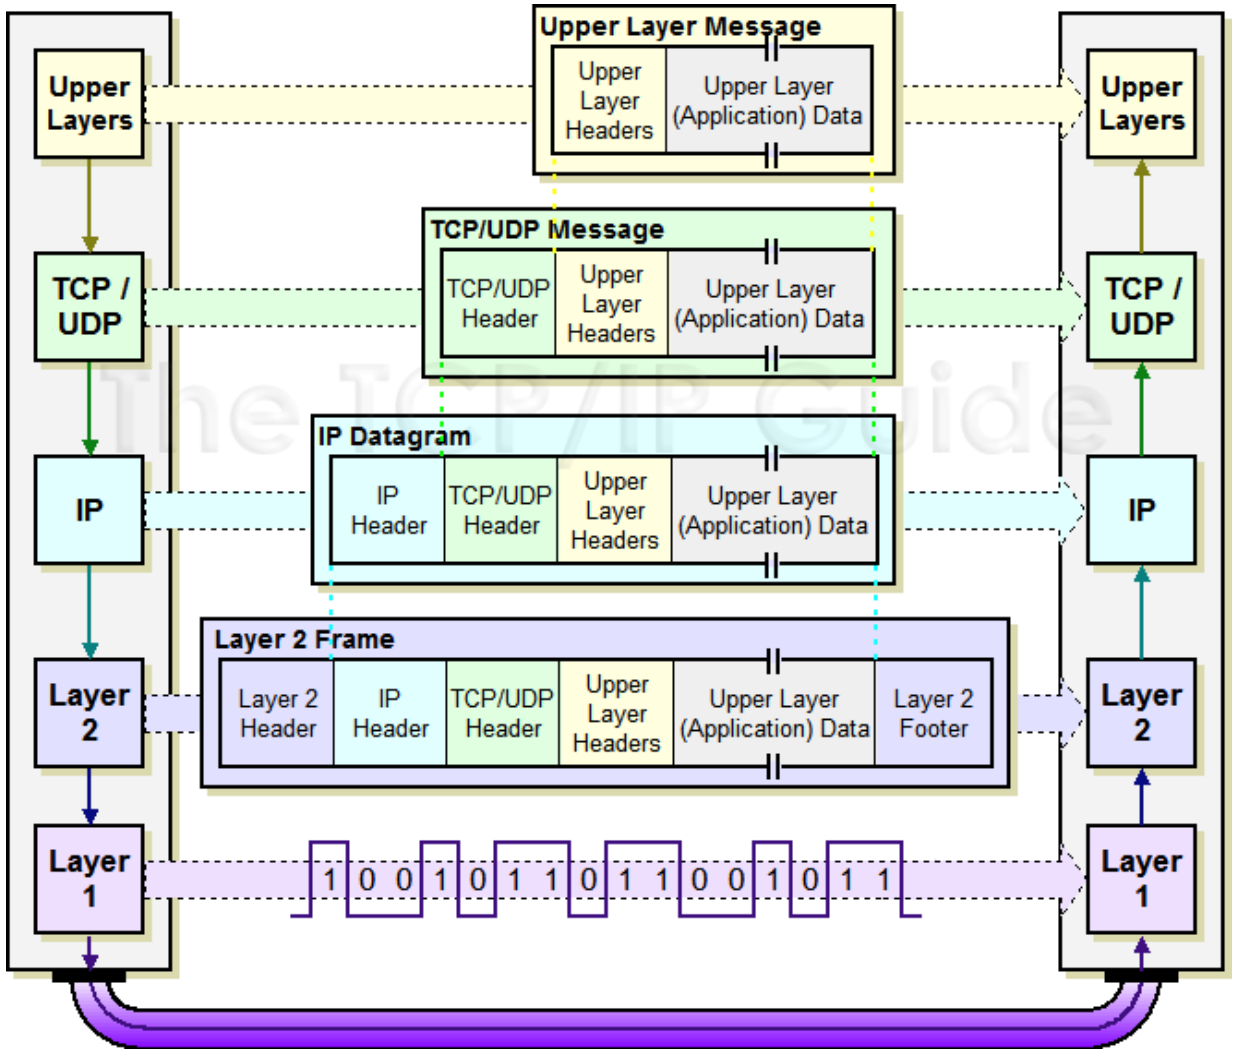

Πανεπιστήμιο Κύπρου

- Host and port scanning
  - Used by hackers during active information gathering
  - They will try to hide the requests within the normal variance of network traffic
- Port sweep
  - One attacker looks for a specific service on many machines
  - More sensitive than port scanning: likely that service is vulnerable (0-day, unpatched)
- Malicious traffic
  - Targeted attacks via network connections
  - Exploitation of networking stack implementation
  - Exfiltration of data
- Distributed Denial of Service (DDoS)
  - Flood a target with network request
  - Typically attacker uses a botnet to diversify source of attack

- Firewalls
  - Examples: iptables (Linux), pf (OpenBSD, OSX)
  - Beyond packet filtering, firewalls can keep state, inspect application-layer packets
  - Can protect individual machine from port scanning, malicious traffic
  - Needs kernel-level operations: critical to defend from compromise
  - No help against DDoS or port sweep
- Intrusion Detection Systems
  - Example: Snort
  - Dedicated hardware that inspects network traffic (on- or off-path)
    - Signature based
    - Anomaly detection based (also machine learning techniques)
  - Can detect host and port scanning, port sweep
  - Can filter malicious traffic but must be fast (hence miss acks)
  - Can rate-limit connections to mitigate DDOS
- But the source IP is easy to spoof: risk of blocking too many IPs
- Variants: egress filters, IPS, etc..

- Causes a service disruption and takes one of two forms:
  - **Receiver Jamming** where a malicious user (called the jammer) emits a signal (called the jamming signal) that interferes with the legitimate signals that are received at the receiver side. The interference degrades the quality of the received signal causing data errors.
  - **Transmitter Jamming** where the jammer in this attack sends a jamming signal that prevents near by senders from transmitting their packets as they sense the wireless channel to be busy and back off waiting for the channel to become idle.

- Constant Jamming
  - Attacker continuously transmits a random jamming signal
    - Can be detected easily
    - requires lots of energy
- Deceptive Jamming:
  - jammer conceals its malicious behavior by transmitting legitimate signal patterns that follow the structure of the MAC protocol rather than sending random bits.
- Reactive Jamming
  - suitable for the case when the jamming device has a limited power budget. The jammer in that case listens to the medium and transmits a jamming signal only after it senses that a legitimate signal is being transmitted in the medium.
- Random Jamming
  - The jammer alternates between sending a jamming signal and remaining idle for random periods of time in order to hide the malicious activity.

- **Frequency Hopping**
  - sender and receiver switch from one frequency to another in order to escape from any possible jamming signal
  - Switching from a frequency to another is based on a generated random sequence that is known only for the sender and receiver.
- **Spread Spectrum**
  - This technique uses a hopping sequence that converts the narrow band signal into a signal with a very wide band, which makes it harder for malicious users to detect or jam the resulting signal. This technique is also very efficient when the transmitted data are protected by an error-correction technique as it allows the reconstruction of the original signal even if few bits of the transmitted data were jammed by the attacker.
- **Directional Antennas**
  - The use of directional antennas can mitigate jamming attacks from being successful as the sender and receiver antennas will have less sensitivity to the noise coming from the random directions that are different from the direction that connects the sender and the receiver
- **Jamming Detection**
  - The receiver can detect that it is a victim of a jamming attack by collecting features such as the received signal strength (RSS) and the ratio of corrupted received packets. Advanced machine learning technique can then be used to differentiate jamming attacks from the degradation caused by the poor quality of the channel due to normal changes in the wireless link

- Exploits the battery lifetime where a malicious user misbehaves in a way that makes devices consume extra amounts of power so that they run out of battery earlier thereby causing a service disruption.
- The damage caused by this attack is usually measured by the amount of extra energy that objects consume compared to the normal case when no malicious behavior exists.

- Denial of Sleep
  - Duty-cycling in data link layer protocols was proposed to reduce the power consumption of devices
  - An adversary can prevent objects from switching to sleep by simply sending control signals that change their duty-cycles keeping them active for longer durations.
- Flooding Attack
  - The adversary can flood other nodes with dummy packets and request them to deliver those packets to the device, wasting energy receiving and transmitting those dummy packets.
- Carrousel Attack
  - This attack targets the network layer and can be launched if the routing protocol supports source routing, where the object generating the packets can specify the whole routing path of the packets it wishes to send to the fog device. The adversary in that case specifies routing paths that include loops where the same packet gets routed back and fourth among the other objects wasting their power.
- Stretch Attack: if the routing protocol supports source routing, then a malicious object can send data through very long paths
- The adversary can further amplify the amount of wasted energy by combining flooding attack with carrousel attack and stretch attack.

- This attack takes place in the case when the object can't send its generated packets directly through the network but must rely on other devices that lie along the path toward the destination device to deliver those packets.
- A malicious device in this attack does not forward a portion of the packets that it receives.
  - A special case of this attack is the blackhole attack where the attacker drops the entire set of packets that it receives from the neighboring objects.
- The best way to prevent packet drops from taking place for sensitive IoT applications is to increase the transmission capability of the objects so that they can reach the destination directly without the need for help from intermediate objects.
- Path redundancy is one of the candidate solutions, where each device forwards each generated packet to multiple neighbors.
  - The main limitation of this mitigation technique is that it has a high energy overhead as it increases significantly the traffic.

- A malicious object claims that it has the shortest path to the destination device
- Now all the packets that are originating from the neighboring nodes pass by this malicious node. This gives the malicious node the ability to look at the content of all the forwarded packets if data is sent with no encryption.
- Furthermore, the malicious object can drop some or all of the received packets as we explained previously in the selective-forwarding attack
- Techniques to detect and isolate the malicious device are based on the idea of collecting information from the different nodes reporting their neighboring devices along with the distance to reach those objects.

- To prevent the hosting machine from being compromised
  - Employing good security measure to ensure:
    - proper functioning of the OS
    - adequate protect the physical machine from unauthorised access
    - correct configuration of the system
- To prevent unauthorised access to IoT software modules and traces
  - Ensuring high quality control on the developed software
  - Monitoring for abnormal operation

- At the cloud each IoT application is dedicated one or multiple virtual machines (VMs) where each VM is assigned to one of the hosting servers in the cloud data center
- VM gets allocated certain amount of CPU and memory resources in order to perform certain computing tasks and store data
- Each IoT application hosted on a VM has its own operating system (OS).
- The hypervisor (sometimes also called the virtual machine manager) monitors those running VMs and manages how these VMs share the server's hardware.

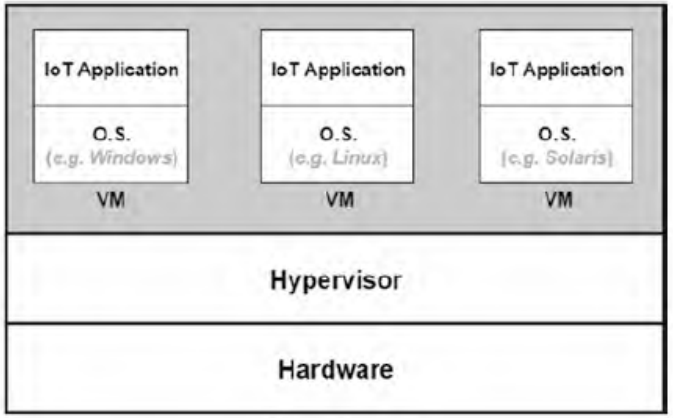

- Although there is a logical separation among the VMs running on the same server, there are still some hardware components that are shared among those VMs such as the cache.
- This opens opportunities for data leakage across the VMs that reside on the same server. Three steps are followed by the attacker in order to leak information from a target VM. These three steps are explained next:
- locate where the target VM resides
    - A cloud data center is typically divided into multiple management units called clusters, where each cluster is located in a certain geographical location and is made up of thousands of servers.
    - Each cluster is divided into multiple zones (sometimes called “pods”) where each zone consists of a large number of servers.
    - Although clients have the choice to specify in which cluster their VM resides (external IP address), they don’t have control on selecting the zone or the server within the zone where their VM will reside as this decision is made based on the cloud provider’s scheduling algorithm which is not released publicly
- Malicious VM Placement
    - place a malicious VM on the same server where the target VM reside by performing traceroute queries
    - Multiple hops indicate different zones
    - Release the rented VM and requests a new one until tracerouter indicates 1-hop distance
- Cross-VM Data Leakage
    - try to learn information about the target VM by exploiting the fact that although VMs are separated logically, thanks to virtualization, they still share certain parts of the server’s hardware such as the instruction cache and the data cache

- Humans introduce further weaknesses
  - Social engineering attacks
  - Do not break defences, but find a way around them
  - Weak passwords
- Insider threats
  - Whistleblowers: Bradley Manning, Edward Snowden
  - Vengeful ex-employees, spies, thieves, etc.
  - Act from a position of privilege: have accounts, know systems and procedures
- Coercion
  - Forced revelation of credentials
  - Lost or stolen devices
- To prevent unauthorised user access
  - Employing strong authentication, such as considering multifactor (what you know, what you have, what you are)
  - Employing good security practice

- A mechanism to ensure authenticity of all modules
- Robust authentication procedure
- Robust trust evaluation
- Software and system certification
- A mechanism to ensure confidentiality and integrity of data
- Strong encryption scheme
- A mechanism to ensure proper behaviour of all modules
- Appropriate monitoring/detection schemes

- Authentication procedure
  - Need complicated handshakes to perform (consume power)
  - Need the client to provide some valid information which can be compromised (eg. password)
- Trust evaluation
- Trust evaluation may rely on recommendations, but they need extensive communications (consume power)
- Encryption scheme
  - Stronger encryption schemes use more power while weaker ones use less (which to use?)
- Monitoring & abnormality detection schemes
  - Each abnormal activity requires a different method to detect (each implementation adds additional load to the IoT system)