# Openstack compliance with GDPR

25th of May 2018 is closer than you think!

FOSDEM 2018

CANONICAL

Vincenzo Di Somma
CISSP
vincenzo.di.somma@canonical.com
@vds

# Agenda

- Introduction
- Why should we care?
- What should we do?

# Introduction to GDPR

1

# What is GDPR

The General Data Protection Regulation (GDPR) (Regulation (EU) 2016/679) is:

- a regulation by which the European Parliament, the Council of the European Union and the European Commission
- intended to strengthen and unify data protection for all individuals within the European Union (EU)

GDPR entered into force in May 2016 and will be **applicable as of 25th May 2018**.

# What is GDPR about?

GDPR is about personal data.

- Specifically about EU citizens or residents personal data
- Not related to where those data are stored or where the company is located or incorporated
- GDPR is aimed at giving back control of personal data to citizens and residents.

# Personal Data

Personal data is defined as:

Any information relating to an identified or identifiable natural person (Data Subject).

# Data Subject

A 'Data Subject' or 'identifiable natural person' is defined as:

one who can be identified, directly or indirectly, in particular by reference to an identifier such as:

- a name
- an identification number
- location data
- an online identifier

or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.

# Processing

Processing is defined as:

Any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as:

- collection
- recording
- organisation
- structuring
- storage
- adaptation or alteration
- retrieval
- consultation
- use
- disclosure by transmission
- dissemination or otherwise making available
- alignment or combination
- Restriction
- erasure or destruction

# Data Controller

A data controller is defined as:

The natural or legal person, public authority, agency or other body which, alone or jointly with others, determines the purposes and means of the processing of personal data.

# Data Processor

A data processor is defined as:

> The entity that processes data on behalf of the Data Controller.

# Why should we care?

2

# Art. 32 GDPR Security of Processing

*“Taking into account the state of the art, the costs of implementation and the nature, scope, context and purposes of processing as well as the risk of varying likelihood and severity for the rights and freedoms of natural persons, the controller and the processor shall implement appropriate technical and organisational measures to ensure a level of security appropriate to the risk, including inter alia as appropriate.”*

# Fines

Maximum fine for non-compliance is either:

- €20 Million
- 4% of an organisation's worldwide annual turnover.

**The greater of the two**, **per violation**!

How much is that?

For Google, a fine could be up to $3.5 Billion!

GDPR is **applicable** from the **25 of May 2018**

# Application vs Infrastructure

- GDPR is about personal data and personal data processing
- Data processing happens mostly at application level not at infrastructure level
- Nevertheless, infrastructure is where data are stored, consider things like:
  - Data at rest
  - Data in transit
  - Backups
  - Logs
  - ...

# Personal Data Breach

This is defined as:

A breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, personal data transmitted, stored or otherwise processed.

# Breach != Fine

What happens if I discover a breach?

- Reporting personal data breaches within 72 hours from discovery is **mandatory**
  - ➢ To the supervisory Authority (Art. 33).
  - ➢ To the Data Subject (art.34)

- Being breached does not automatically imply a fine.
- But breach will likely imply an audit.

# What should we do?

3

# Company policies

Companies dealing with personal data must put in place GDPR compliance policies and guidelines.

Our job is to make sure our Openstack deployments fit perfectly into the company policies and guidelines for GDPR compliance.

# Data Protection Officer

The role of the Data Protection Officer is

- to be an expert on data privacy
- who works independently
- to ensure that an entity is adhering to the policies and procedures set forth in the GDPR

# Data protection by design and by default

Privacy and data protection:

- must be a key consideration in the early stages of any project, and throughout its lifecycle
- The concept of ‘Privacy by Design’ already exists, it has now been given specific recognition, and is linked to enforcement

Or, if you prefer, GDPR delineates ‘Privacy by Design’ as a ‘legal obligation’.

# Best Practices

In open source projects, best practice is often defined by the community and easily available.

- Openstack Security Guide:

  https://docs.openstack.org/security-guide/

# Release Series and End of Life

Avoid unmaintained releases!

- OpenStack is developed and released around 6-month cycles
- After the initial release, additional stable point releases will be released in each release series
- End of Life is after around 12 months from release date
- https://releases.openstack.org/

Soon Ocata will reach end of life: 2018-02-26.

Plan ahead for the migration.

# vulnerability:managed

## Barbican (Key Manager service)

**Home Page:** https://wiki.openstack.org/wiki/Barbican
**PTL:** Dave McCowan (dave-mccowan)
**IRC Channel:** #openstack-barbican
**Service:** Key Manager service

### Mission

To produce a secret storage and generation system capable of providing key management for services wishing to enable encryption features.

### Team-based tags

- team:diverse-affiliation

### Deliverables

#### barbican

**Repositories:** openstack/barbican
**Tags:**
- vulnerability:managed
- assert:follows-standard-deprecation
- assert:supports-upgrade
- stable:follows-policy

#### barbican-specs

**Repositories:** openstack/barbican-specs

# Openstack Projects 'vulnerability:managed'

This tag is part of the vulnerability-classification system for vulnerability reporting and tracking across project deliverables.

**vulnerability:managed**

- indicates that a deliverable vulnerability report reception and disclosure are handled directly by the OpenStack Vulnerability Management team (VMT)

# Vulnerability Managed Openstack Projects

**barbican** (Barbican (Key Manager service))
**castellan-ui** (Barbican (Key Manager service))
**python-barbicanclient** (Barbican (Key Manager service))
**cinder** (Cinder (Block Storage service))
**python-cinderclient** (Cinder(Block Storage service))
**glance** (Glance (Image service))
**glance-store** (Glance (Image service))
**python-glanceclient** (Glance (Image service))
**heat** (Heat (Orchestration service))
**python-heatclient** (Heat (Orchestration service))
**horizon** (Horizon (Dashboard))
**keystone** (Keystone (Identity service))
**python-keystoneclient** (Keystone (Identity service))
**neutron** (Neutron (Networking service))
**neutron-lib** (Neutron (Networking service))
**python-neutronclient** (Neutron (Networking service))
**nova** (Nova (Compute service))
**python-novaclient** (Nova (Compute service))
**castellan** (Oslo (Common libraries))
**oslo.config** (Oslo (Common libraries))
**python-saharaclient** (Sahara (Data Processing service))
**sahara** (Sahara (Data Processing service))
**sahara-dashboard** (Sahara (Data Processing service))
**sahara-extra** (Sahara (Data Processing service))
**sahara-image-elements** (Sahara (Data Processing service))
**python-swiftclient** (Swift (Object Storage service))
**swift** (Swift (Object Storage service))
**python-troveclient** (Trove (Database service))
**trove** (Trove (Database service))

# Logging And Monitoring

- Accountability
- Be aware of the status of the system
- Timely identify breaches
- Reporting breaches is mandatory
- Support Forensic

# Pseudonymisation of the Logs

- The processing of personal data in such a manner that the personal data can no longer be attributed to a specific data subject without the use of additional information, provided that such additional information is
  - kept separately, and
  - subject to technical and organisational measures to ensure that the personal data are not attributed to an identified or identifiable natural person.

# Backup, Decommission, Data Wiping and Right to be Forgotten

- Remove personal data while backup restore.
- Be careful when decommission hardware.

# Automation is the key to compliance

- Openstack is a complicated piece of software.
- You are going to have tons of configurations.
- Your security policies will be implemented in your configurations.
- It’s best practice.
- Plenty of tools, Juju + MAAS + Charms, Openstack/Ansible, ...




# Openstack Networks and Segregation

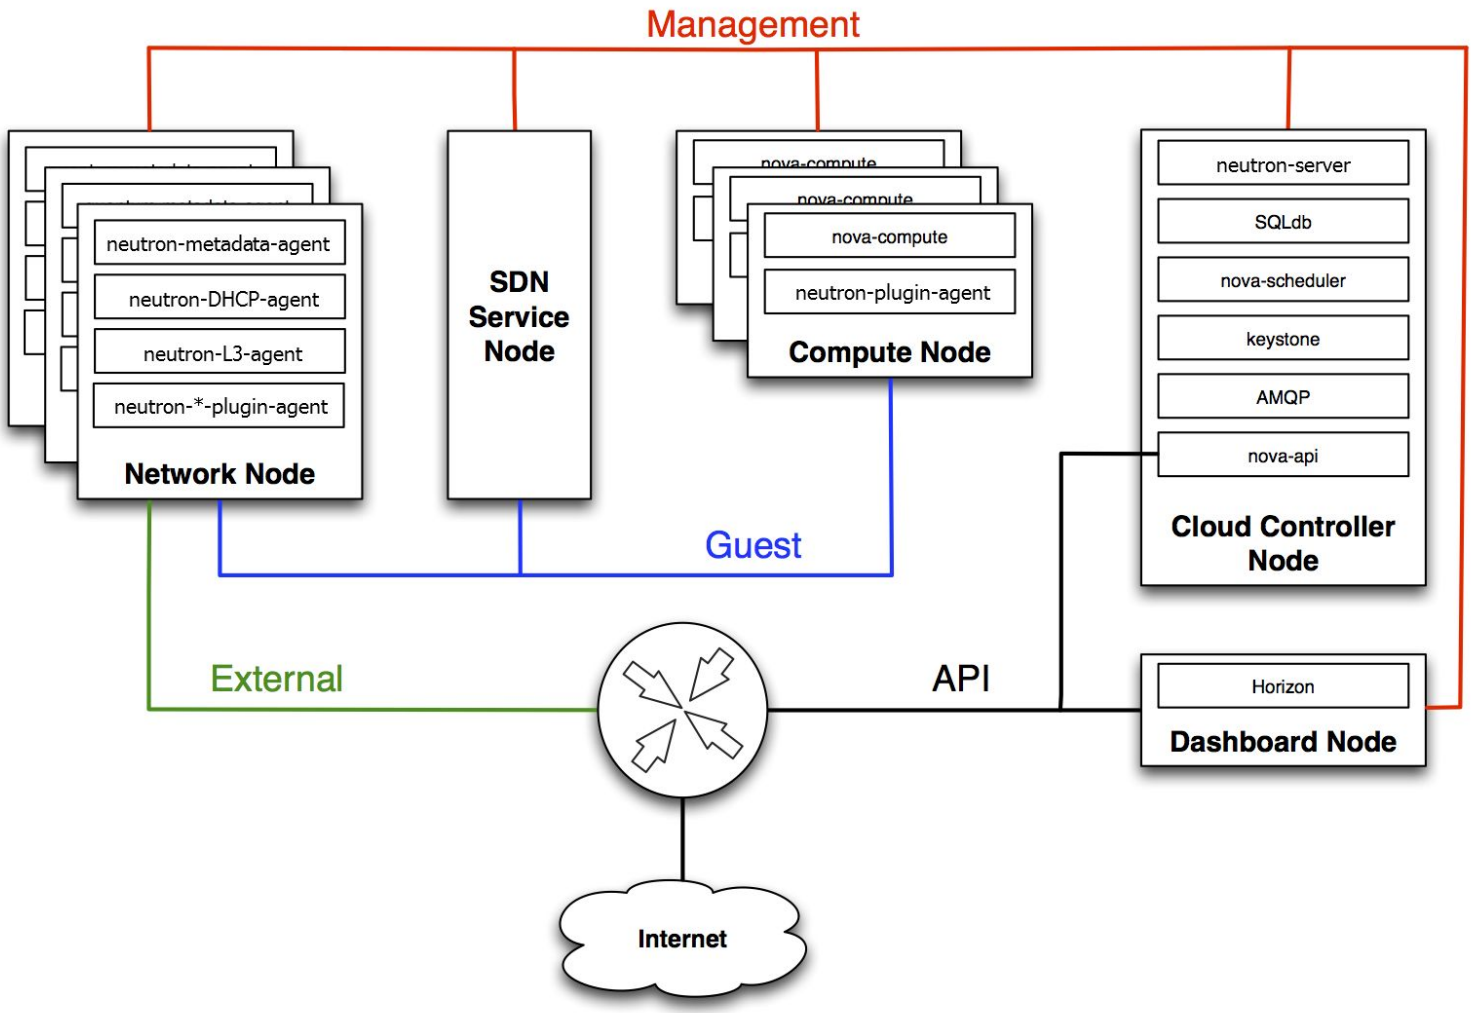

# Encryption of data in transit - TLS

- Ideally use SSL/TLS on both public networks and management networks in TLS proxies and HTTP services
- If it’s not doable carefully identify the threats OpenStack and use SSL/TLS accordingly

# Encryption of data in transit

- Publicly facing services, deal with credentials
- The attacker can then use these valid credentials to perform malicious operations
- All real deployments should be using SSL/TLS to protect publicly facing services

# Encryption of data in transit

- Services that are deployed on management networks, are subject to internal attacks
- Unauthorized users might gains access to the internal network by exploiting a misconfiguration or software vulnerability
- Using SSL/TLS on the management network can minimize the damage that an inside attacker can cause

# Encryption of data at rest - Storage Encryption

- Volume encryption
- Ephemeral disk encryption
- Object Storage objects

# Intrusion Detection and Prevention Systems

- A leading cause of incidents that compromise personal data remains external breaches, caused by attackers gaining unauthorized access to a network
- IDS-IDP are the best safeguards against these risks
- These tools are able to monitor any traffic coming into and moving within a network and alert businesses if any suspicious activity is detected

# Malware Protection

Malware protection on Linux???

- ClamAV
- chkrootkit and rkhunter

# Define robust  policies for handling PKIs.

Public Key Infrastructure: framework for creating a secure method for exchanging information based on public key cryptography.

Need to define:

- policy on how to deal with root CA
- a robust process to checkout certificates
- how chain certificates will be used

PKIs have been around for many years, and still have a lot of issues, unfortunately, auditors are skilled in finding issues of PKIs.

This is where automation will help.

# Security Testing

Pen testing your Openstack instance should happen not only after the deployment but periodically.

Syntribos is an open source automated API security testing tool that is maintained by members of the Openstack Project.

- https://github.com/openstack/syntribos

# Train Your Employees

- Staff need to understand the GDPR
- Training must be relevant
- Simulations are helpful in ensuring understanding
- Staff should be able to identify breaches and "red flag" situations

# Security Incident Response Team

- What happened?
- Why did it happened?
- Who was affected?
- How to prevent it from happening again?

# Spectre and Meltdown

- Spectre breaks the isolation between different applications, CVE-2017-5753 and CVE-2017-5715.
- Meltdown exploits side effects of out-of-order execution to read arbitrary kernel-memory locations, CVE-2017-5754.

An attacker in a hostile VM running on an unpatched host kernel could access data in other VMs running on the same host.

**Keep your kernels patched!**

# Thank you!

Vincenzo Di Somma
CISSP
vincenzo.di.somma@canonical.com
@vds