Red Hat UK Forum 2018 - Ideas Worth Spreading

Use of open source technologies to deliver modern public health services

Francesco Giannoccaro

Public Health England

Protecting and improving the nation's health

Public Health England

Summary

Benefits and challenges of using open source HPC and Cloud technologies to deliver modern public health services and to facilitate data sharing and accessibility

Public Health England

About PHE

- Public Health England (PHE) is an executive agency of the Department of Health in the UK. We provide government, local government, the NHS, Parliament, industry and the public with evidence-based professional, scientific and delivery expertise and support.

- PHE mission is to protect and improve the nation's health and wellbeing, and reduce health inequalities. We do this through world-leading science, knowledge and intelligence, advocacy, partnerships and the delivery of specialist public health services.

- Public Health England was established in 2013 to bring together public health specialists from more than 70 organisations into a single public health service. We employee about 5,500 staff, mostly scientists, researchers and public health professionals.

Public Health England

Wide range of public health services

PHE deliver a wide range of public health services including

- research and scientific publications based on mathematical models such as Spatial Metapopulation Model for transmissible disease (eg Flu/Smallpox), predictive models applied to the Anthrax inference problem to be able to infer: likely size of outbreak, location of source, spatial extent, etc

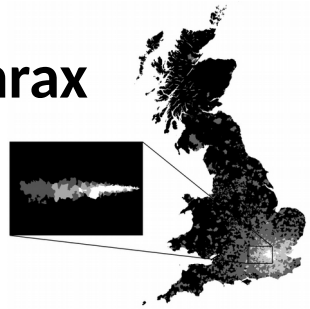

- pathogen genomics service, based on whole genome sequencing, for pathogen typing, surveillance and outbreak investigation. More than 100,000 bacterial and viral genomes have been sequenced since the service launch in 2014 (updated to Q1/2018)

- campaigns such as: Be Clear on Cancer, Act FAST, Stoptober, Change4Life, etc

Public Health England

Challenges of structuring a central ICT department merging more than 70 different organisations

We focused on shaping the ICT strategy to better support our scientific community and designing a technology innovation roadmap to shift from a restricted stack of proprietary technologies into a more Open orchestrated ecosystem

initial technology stack managed by our central ICT department: the main focus was on supporting business as usual requirements

some of the new open source technologies that we have introduced during the redesign and innovation of our infrastructure to support the PHE scientific community

Public Health England

Use of HPC in Public Health England - Background

The **Bioinformatics** Unit

- Processing and analysing DNA for diagnostics and surveillance of infectious diseases. Samples received from patients with unidentified and potentially aggressive pathogens (bacteria and virus) that need urgent identification.
- Mainly **high throughput computing**: many small jobs, lots of CPU and disk I/O

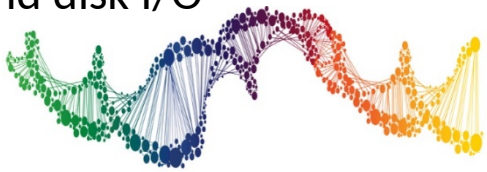

The **Modeling and Economics** Department

- Running multiple real time models and simulations to predict expected pandemic disease dynamics, supporting national vaccination policy and control of antimicrobial resistance;
- **Traditional MPI** HPC: larger jobs, lots of CPU and small amount of disk I/O

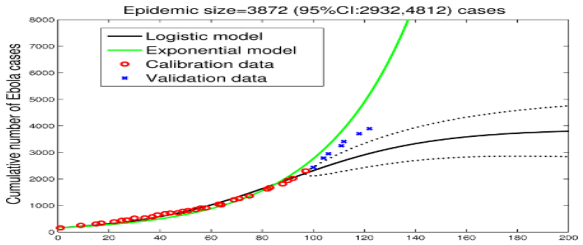

Emergency Response Department

- Running multiple models and simulations to better understand ahead of time the epidemiological, social, behavioral drivers that exacerbate the risks posed by infectious disease threats, including bioterrorism;
- Mainly **traditional MPI** HPC and some GPU: larger jobs, lots of CPU and moderate disk I/O.

Public Health England

The initial HPC systems where fragmented and isolated

High Performance Computing System used by Bioinformatics unit
416 cores, 3.3TB of RAM

Located in Colindale

Linux cluster based on RHEL
- Resource manager: GridEngine
- Provisioning system: xCat

2 x Management server IBM x3650
- 2 x Intel X5450, 32GB of RAM
- 2 x 10 GB Ethernet
- 6 x 72 GB SAS

16 x HP Blade BL460c Gen8 compute nodes
- 2 x Intel E5-2680, 128 GB of RAM
- 2 x 10 Gb Ethernet
- 2 x 900GB 6G SAS 10K

10 x IBM Flex System x240 compute node
- 2 x Intel E5-2650v2, 128GB of RAM
- 1 x CN4022 2-port 10Gb
- 2 x 900GB 10K SAS HDD

High Performance Computing System used by Modeling and Economics department
864 cores, 3.5TB of RAM

Located in Colindale

Linux cluster based on Bull/RHEL
- Resource manager: Slurm
- Provisioning system: Bull

2 x Bull R423-E3 management server
- 2 x Intel E5-2620, 32GB of RAM
- 2 x 500GB SATA3
- 2 x InfiniBand ConnectX-2 QDR IB
- 2 x Gb Ethernet

72 x Bull B500 compute nodes:
- 2 x Intel X5660, 48 GB of RAM
- 1 x 128GB SATA2 SSD
- InfiniBand adapter

High Performance Computing System used by Emergency Response department
432 cores, 1.2TB of RAM

Located in Porton

Linux cluster based on Bull/RHEL
- Resource manager: Slurm
- Provisioning system: Bull

28 x Bull B510 compute nodes:
- 2 x Intel E5-2620, 32GB of RAM
- 1 x 256 GB SSD
- 2 x 1Gb Ethernet
- 1 x InfiniBand QDR

8 x Bull B500 compute nodes:
- 2 x CPU Intel L5530, 24GB of RAM
- 1 x 256 GB SSD
- 2 x 1Gb Ethernet
- 1 x InfiniBand QDR

1 x Bull GPU server:
- 2 x CPU Inte E5620, 20GB of RAM
- 2 x Nvidia K20c GPU cards
- 1 x InfiniBand QDR - 2 x 1Gb Ethernet
- 1TB SAS disk

High Performance Storage tier based on Lustre, current usable capacity 800TB

Public Health England

A technology innovation program based on open source

ICT in 2015/2016 designed and implemented a **technology innovation program** underpinned by the ICT strategy, focused on **supporting the scientific community** within PHE. This work included the following three core projects:

1) Open source virtualization platform (oVirt/RHV) based on **KVM hyper-visor**

- to reducing vendor lock-in constraints on proprietary virtualization technology
- to provide users with a self service portal and open API for server deployment and management

2) HPC Cloud to support the needs for high performance and high throughput computing, based on the open source technology OpenStack.

3) PHE scientific data catalogue: as a prerequisite to enable Big Data analytics capability in 2015 we started deploying **iRODS** an open-source technology to help our scientists to organise, catalogue, aggregate and share large sets of data that are currently fragmented and isolated in silos

Public Health England

oVirt open-source virtualization platform based on Linux/KVM

In order to reduce vendor lock-in and to have a more cost effective virtualization infrastructure **oVirt** was deployed in 2015. This is an open-source enterprise virtualization platform that provides similar features of proprietary platforms such as VMware

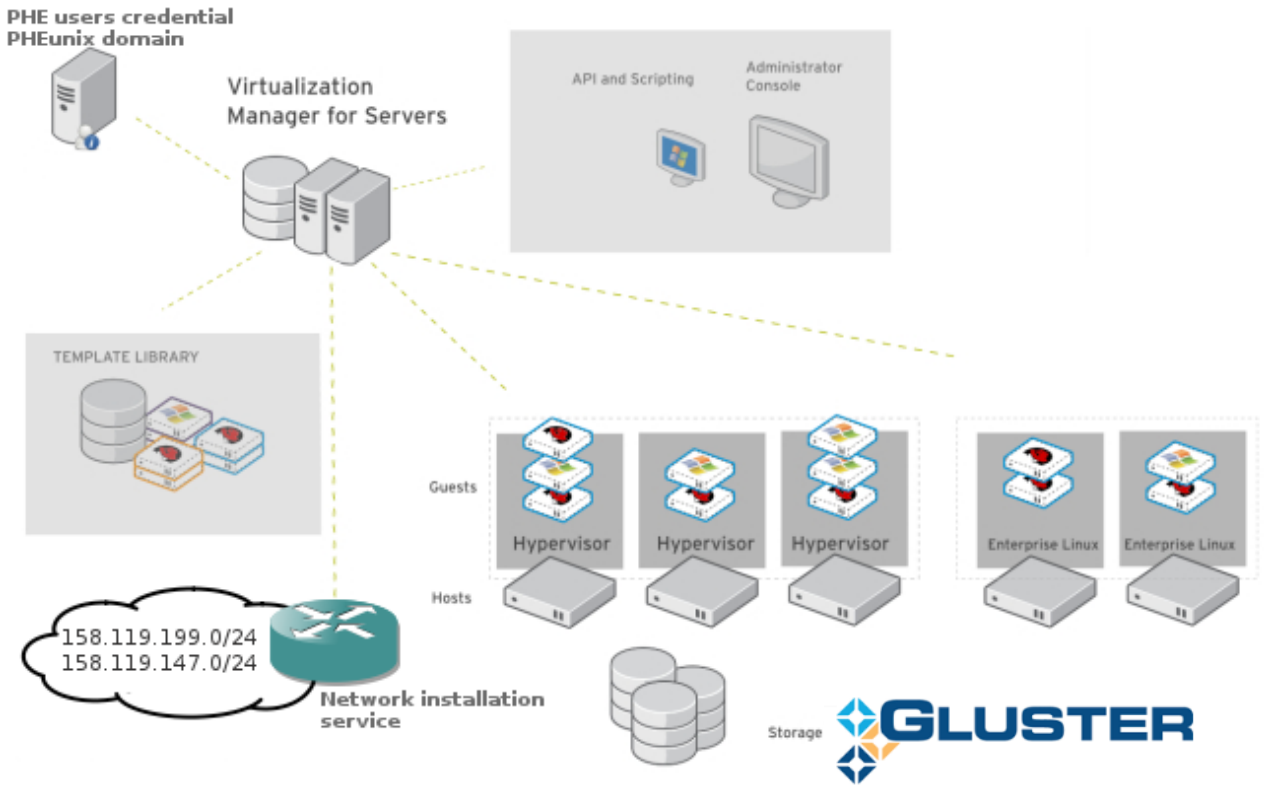

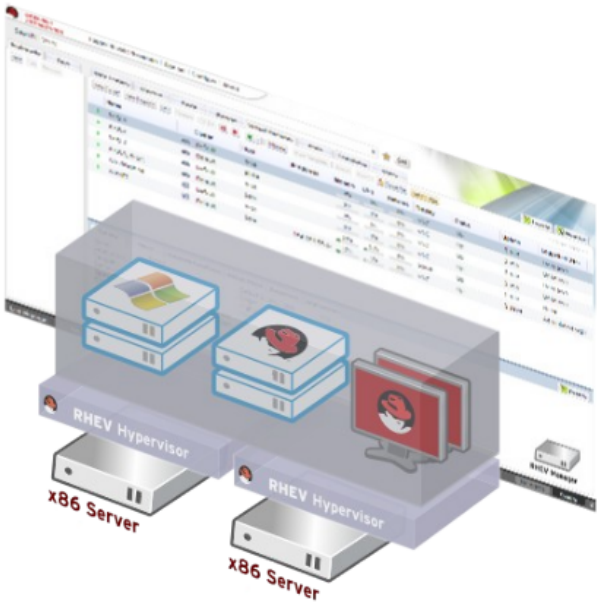

Public Health England

geo-distributed storage system (13 Petabyte) using S3 open standard

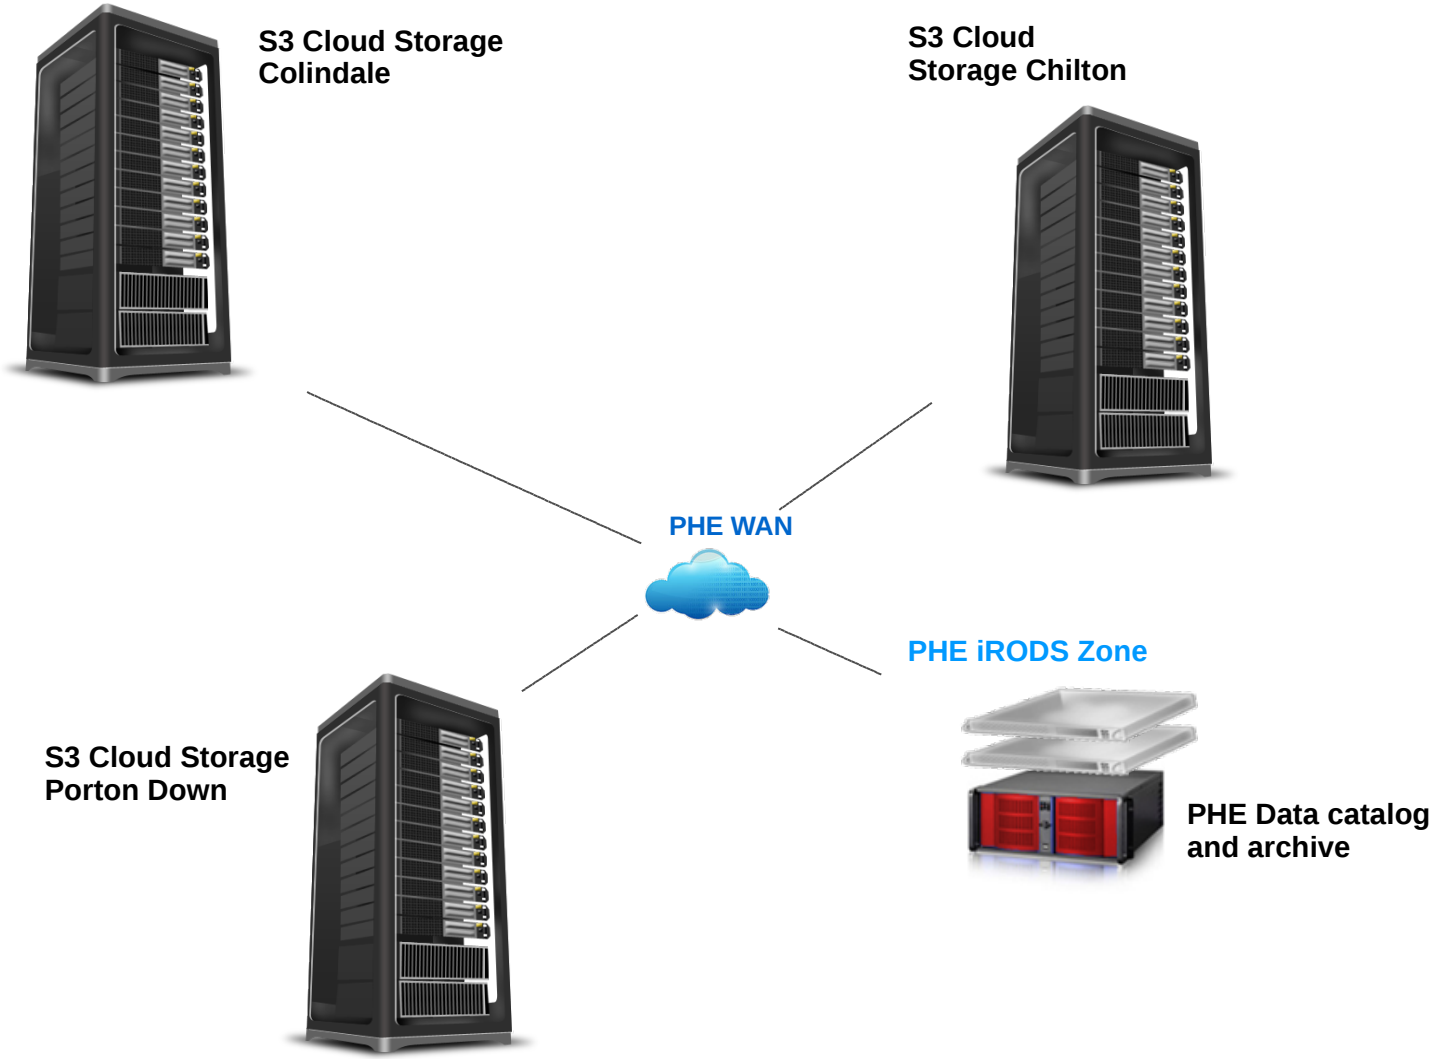

Public Health England

HPC/HTC Cloud project

Goals & objectives

- provide HPC on demand services, offering the ability to expand any of the existing clusters by deploying additional compute resources when needed; provide a cost effective IT infrastructure enabling efficient resources sharing;
- provide the ability to **easily deploy virtual HPC** cluster with improved **orchestration and automation** capability. Using software defined infrastructure technologies to support **elastic allocation of resources**, centralized provisioning and management operations;
- increase resilience and disaster recovery capability by implementing **geographically distributed cloud regions across multiple sites**;

The OpenStack journey

- We wanted to break away from our three standalone HPC clusters, OpenStack would **allow the existing clusters to "burst"** onto the on premise cloud environment. We also wanted to give out users the ability to create instances and create **virtual clusters** (e.g. ElastiCluster, Senlin)
- In **2014** we started a **first** a **proof of concept** (with Red Hat and OCF) running **OpenStack Icehouse**, it was very early days for OpenStack, for us too complex to deploy and manage

OpenStack is a global collaborative project to build the open standard cloud computing platform

www.openstack.org - Q1/2016

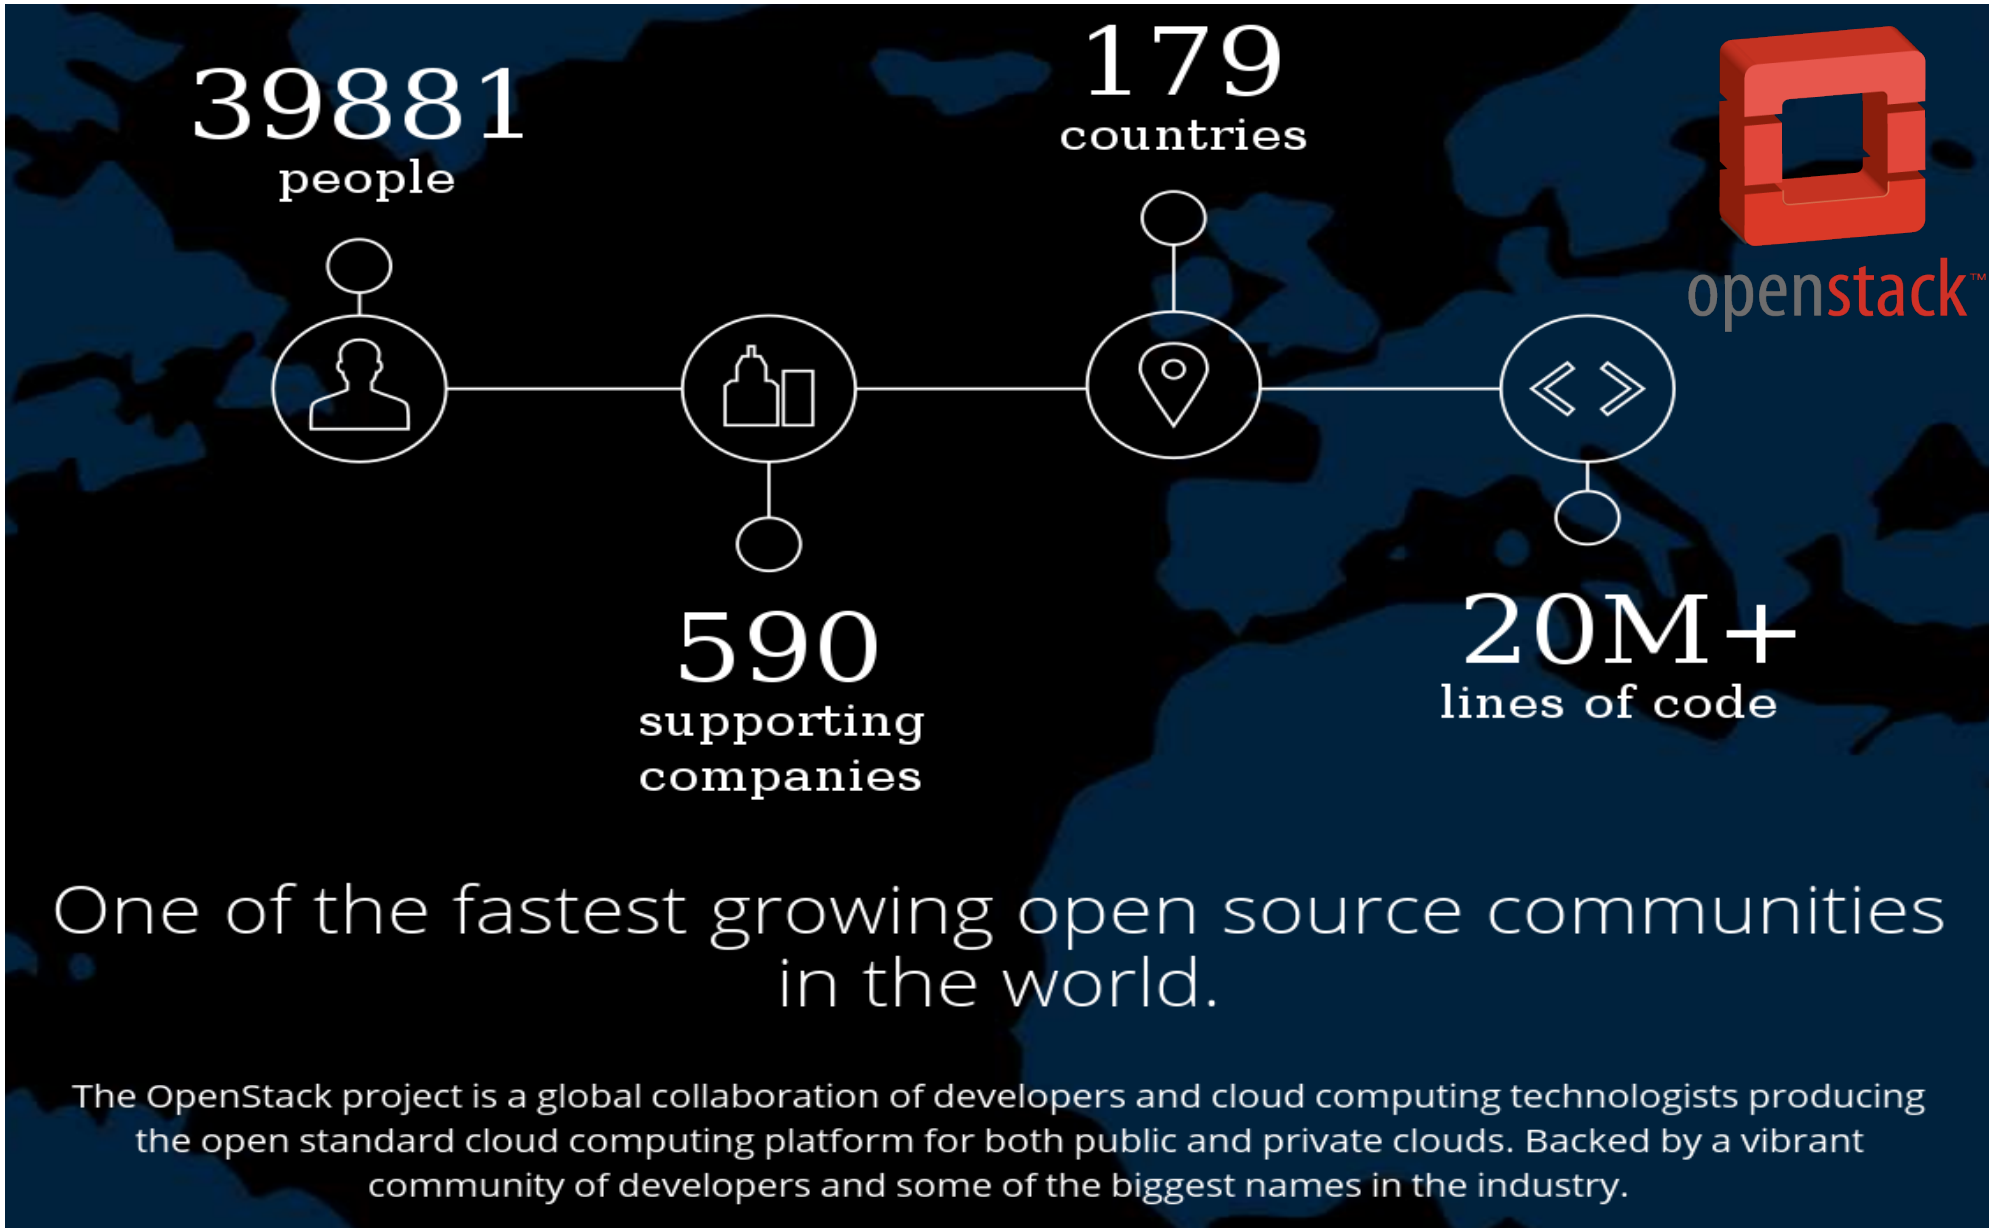

Public Health England

PHE OpenStack deployment across two cloud regions (production 2016)

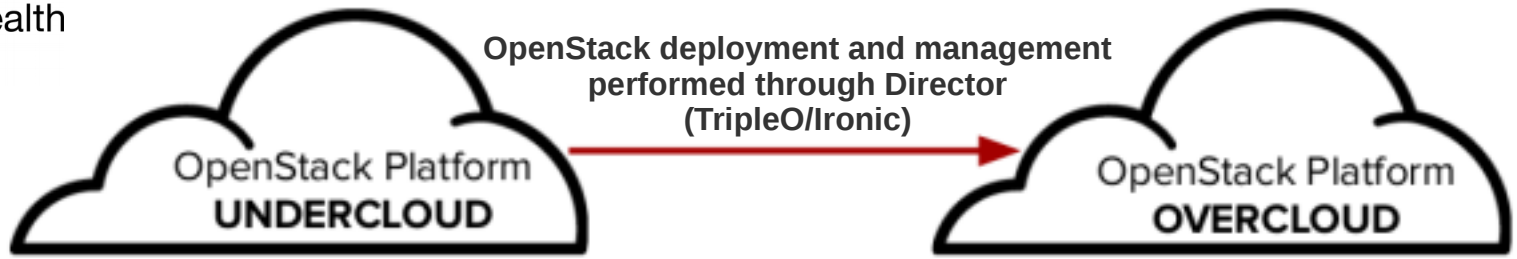

openstack.
Newton (OSP10)

OpenStack Undercloud provisioning servers:

- Lenovo x3550 M5 with 1x Intel E5-2620 v3
- 64 GB of RAM
- 2x 1TB 10K 6Gbps SAS HDD
- 1x ConnectX-3 Pro 2x40GbE/FDR VPI Adapt.

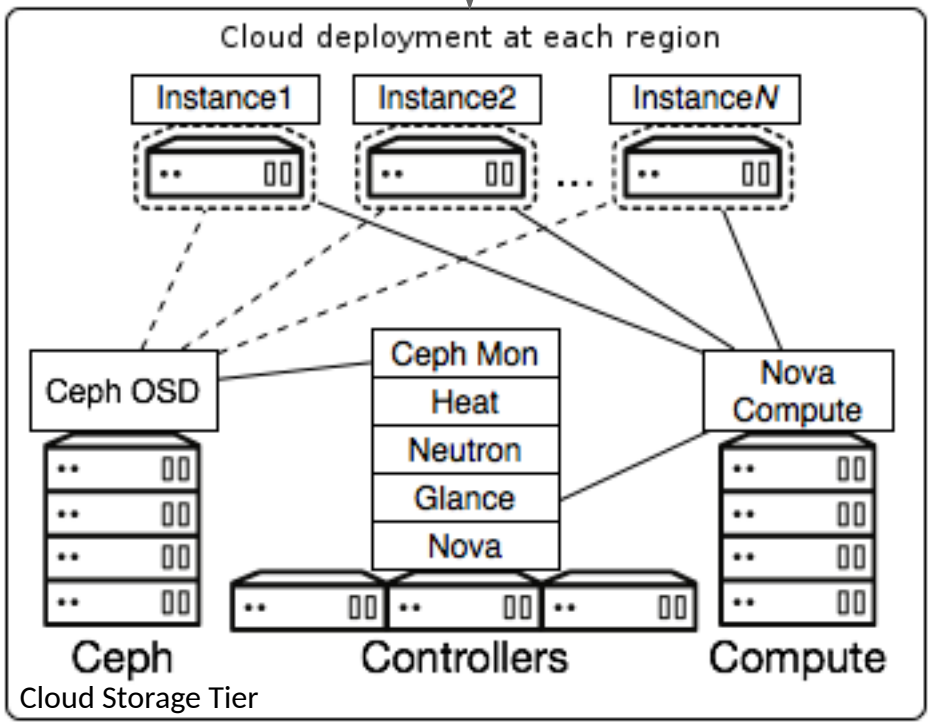

3 x Controller nodes - Lenovo x3550 M5:

- 2x Intel E5-2620 v3, 128 GB of RAM
- 2x 240GB SATA SSD - 2x 480GB STAT SSD
- 1x ConnectX-3 Pro ML2 2x40GbE/FDR VPI Adapter

35 x Compute nodes (*) – Lenovo nx360 M5:

- 2x Intel E5-2640 v3(16cores) , 128 GB of RAM
- 2x 120GB SATA SSD
- 1x ConnectX-3 Pro ML2 2x40GbE/FDR VPI Adapter

1 x **Compute node for large cloud instances** - Lenovo x3950 X6 8U

- 8x Intel E7-8860 v3(128cores) - 1TB of RAM
- 2x 120GB STAT SSD
- 1x ConnectX-3 Pro ML2 2x40GbE/FDR VPI Adapter

2x **GPU nodes** – Lenovo nx360 M5:

- 2x Intel E5-2640 v3, 128 GB of RAM
- 2x 120GB STAT SSD
- 1x ConnectX-3 Pro ML2 2x40GbE/FDR VPI Adapter
- 1x nVidia Tesla K80

Public Health England

PHE 2018/19 cloud development road-map

- Standardise **system provisioning and configuration management** process (through **Ansible AWX**) as well as improving building-images process
- Upgrade **Lustre** to provide **multi-tenancy** oriented integration with OpenStack
- Improve **cross-platform orchestration** (through **ManageIQ/CloudForms**), enabling **centralized governance and management of hybrid/multi clouds**, standardising procedures across all PHE platforms (OpenStack, Vmware, oVirt/RHV)
- **Resource quota enforcement, and accountability for public cloud usage** (ManageIQ)
- Ensure **compliance with security policies**, **configuration auditing** and **change tracking** (through GitLab);
- **Expand use** of **GPU nodes**, and use more extensively OpenStack projects ecosystem (Sahara, Trove)
- Improve reproducibility and **workload scalability and portability** through use of **containers** (Kubernetes, Docker, Singularity)

Public Health England

PHE 2018/19 cloud development road-map: cross platform Orchestration layer to manage hybrid clouds and multi clouds

ManageIQ (CloudForms is a commercially supported version of ManageIQ) is an open source **orchestration layer** that provides a unified and consistent set of management capabilities across multiple virtualization and cloud platforms (OpenStack, AWS, Google Cloud, VMware, Microsoft Azure, etc). It:

- offers IT **infrastructure resources as a service**, in a consumption oriented way, supporting broad **cross-platform end-user service catalogue**;
- allows dynamic service design and **cross-platform infrastructure modelling**;
- has expandable **plugins** allowing connections between private and **commercial clouds** enabling ICT **to manage hybrid and multi cloud** in a fully distributed cloud infrastructure;
- offers a consistent user experience and functionality across all underpinning platforms;
- **reduces vendor lock-in** and supports the organization with the technology life-cycle transition process.

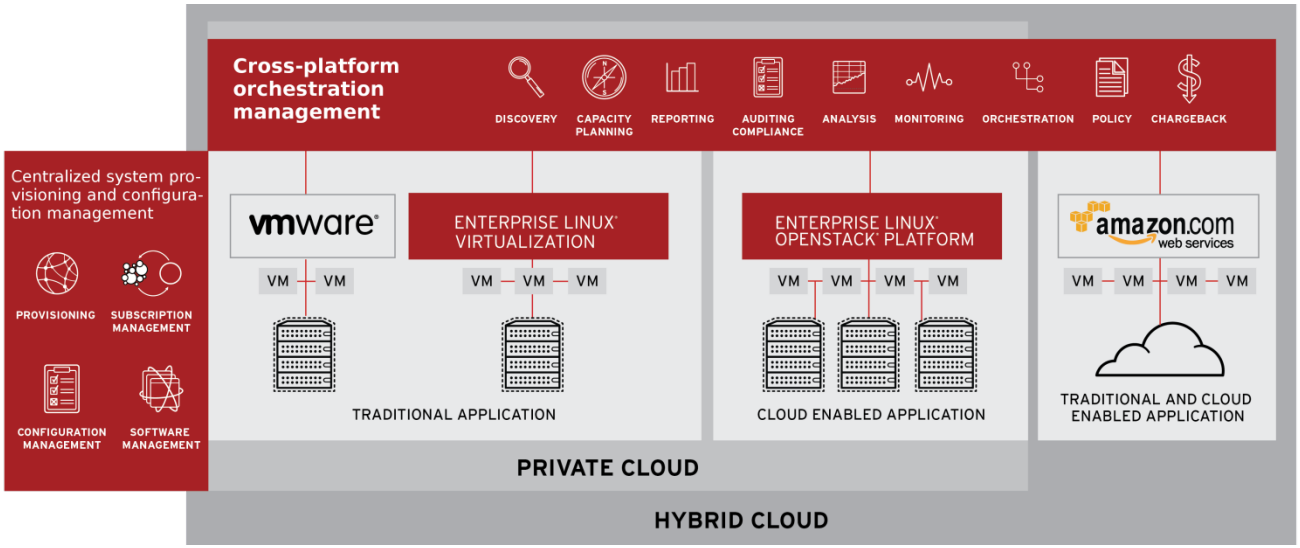

Public Health England

Managing PHE Infrastructure as Code, improving automation and centralized configuration practices

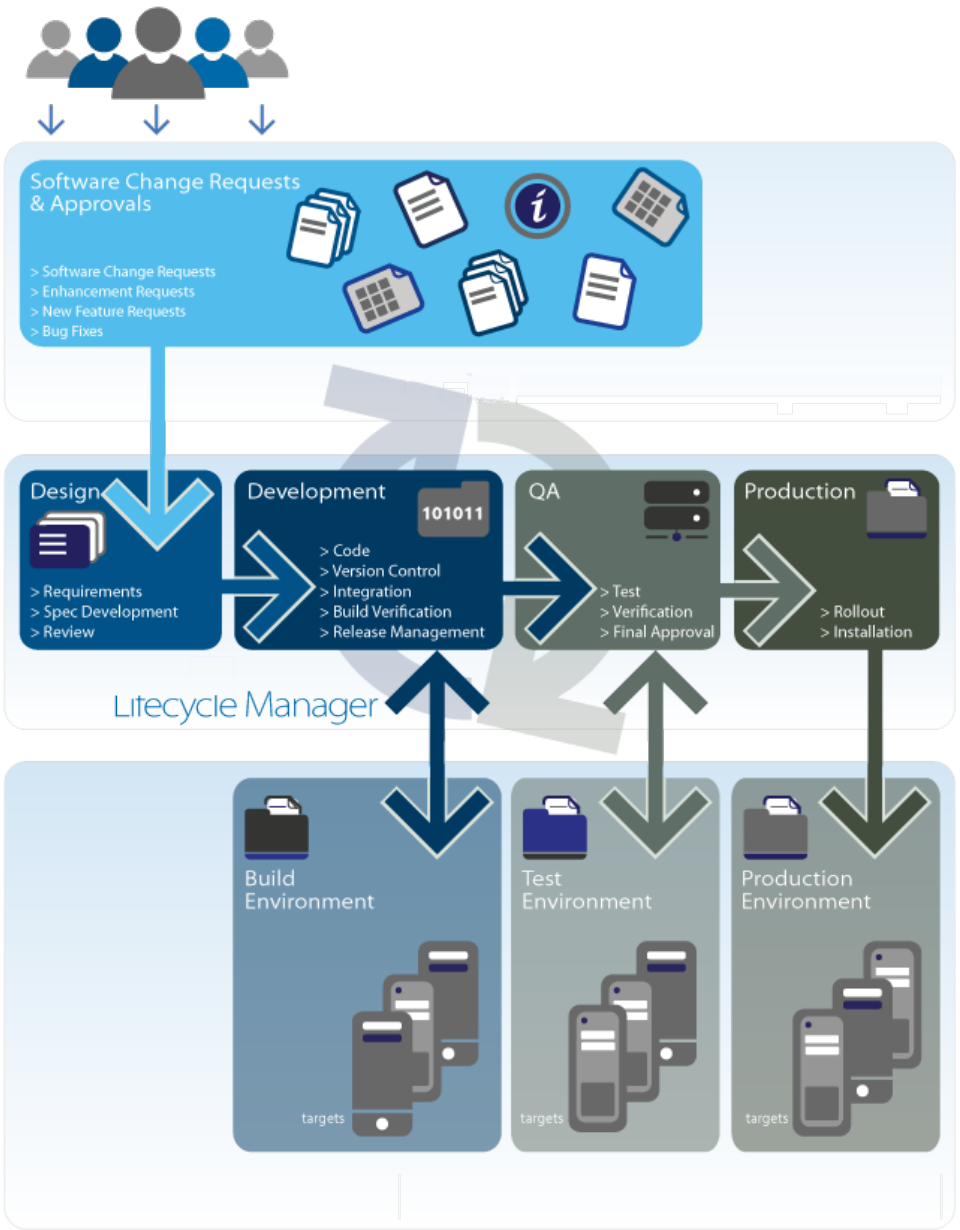

The broad ecosystem of technologies that will be implemented in 2018 (MangeIQ API and Ansible) will accelerate service delivery and reduce operational costs through a **centralized configuration management system** that will provide:

- **automatic provisioning**;
- resource quota enforcement;
- ensure **compliance** and governance;
- configuration auditing and change tracking.

This will ensure that systems will be configured through configuration definition files, as scripts. This will also mean that applying configuration updates with code/scripts will be **fast and consistent**, allowing engineers to provision many servers faster and with less risks of error than any human could do through "type and click" procedures.

Public Health England

Unified cross-platform management with ManageIQ/CloudForms & Ansible

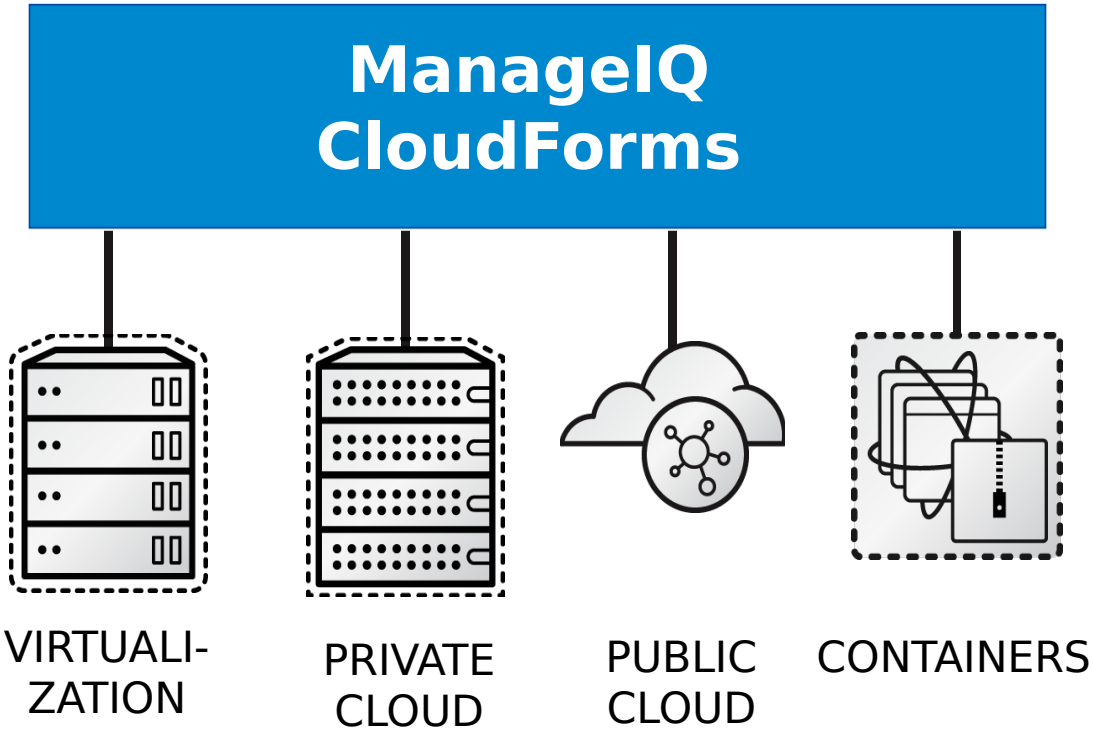

- have one management system for all of our platforms
- have consistent automation & policies that are common across platforms
- facilitate cross-platform infrastructure governance and centralised costs monitoring

Public Health England

Overview of the PHE infrastructure

Orchestration layer to manage on-premise virtualization platforms hybrid and multi cloud distributed infrastructure

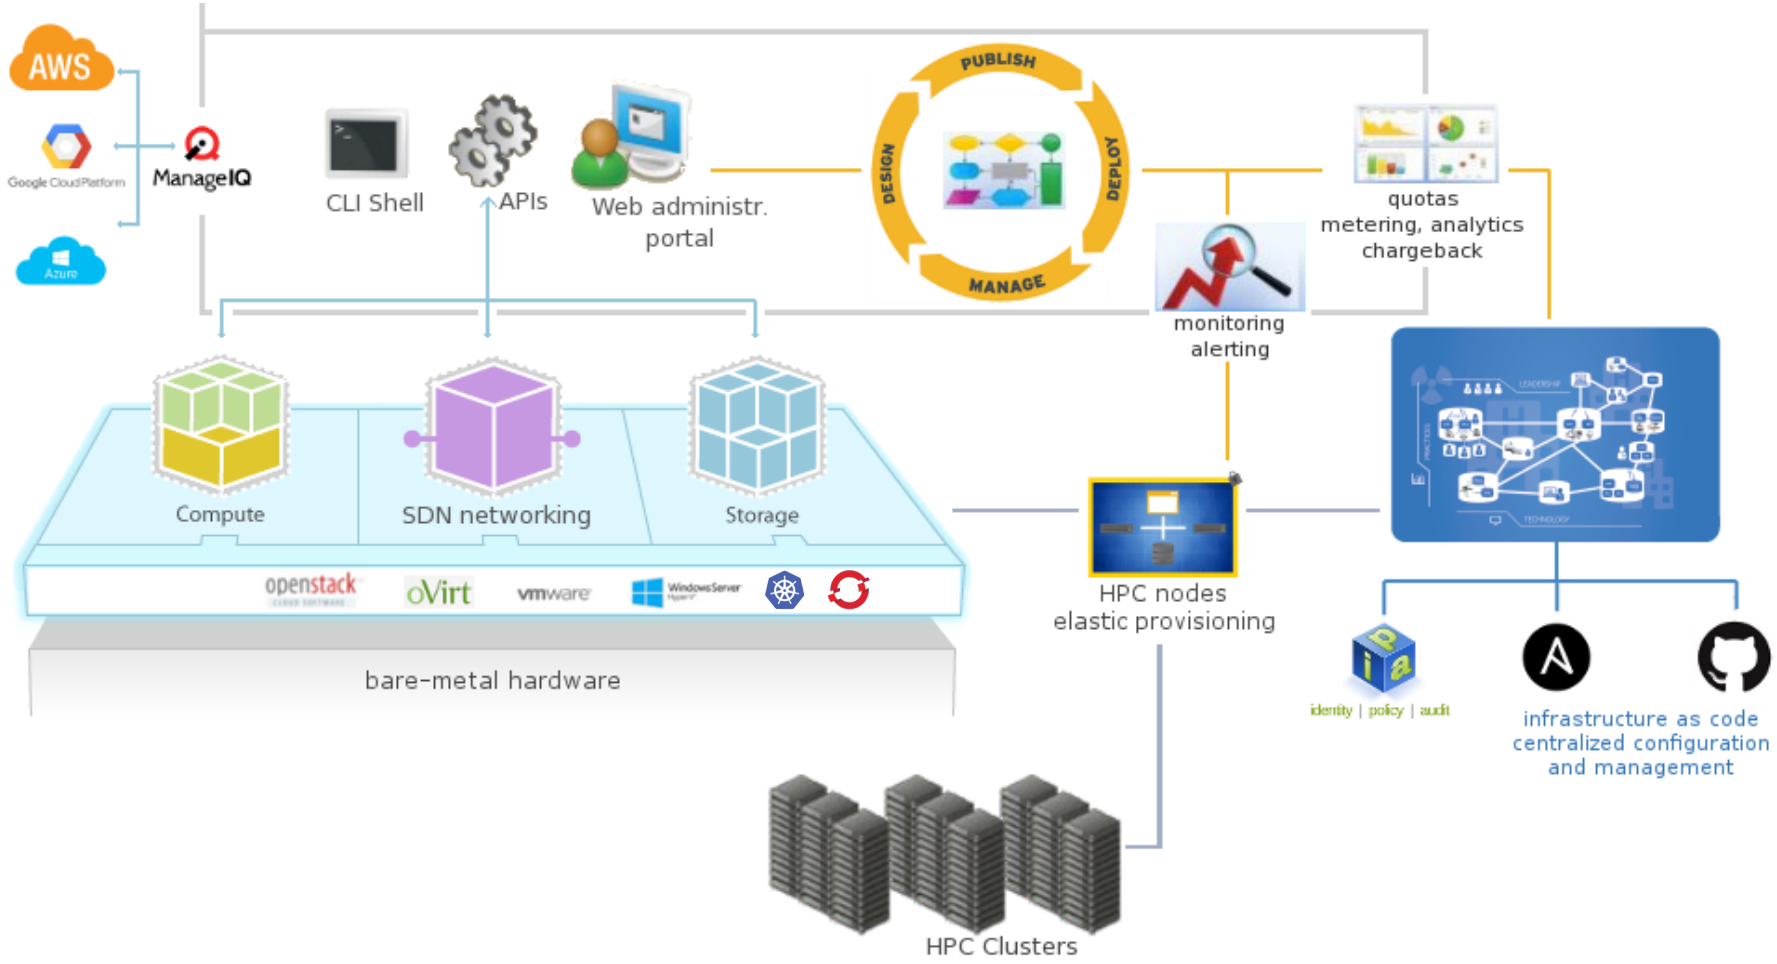

HPC Infrastructure capacity	Cores	Ram	HPS	Archive Storage
Bare-metal HPC clusters	1.5k	6.4TB	400TB	500TB
HPC/HTC Cloud	**3k**	**17TB**	**800TB**	**8 PB**

Public Health England

Acknowledgements & thanks

Team members and key contributions

Public Health England

Thomas Stewart, Tim Cairnes, Anna Rance, Hemanth Manupati, Sam Morris and many other PHE colleagues

Technical partners

OCF redhat DataDirect NETWORKS Mellanox TECHNOLOGIES

A special thanks to the open source community

Keep in touch

francesco.giannoccaro@phe.gov.uk www.linkedin.com/in/giannoccaro