

# General Purpose Servers

Industry Day, Bologna, 16 January 2019

Manuel Viceiro

Team Leader - Servers and Desktops

Platforms and Services Section



# Disclaimer

- All names or descriptions used in this presentation which are trademarks, trade or brand names, or other references to proprietary products are hereby acknowledged as the property of their respective owners. No entry, term or definition in this presentation should be regarded as having any implication as to the validity or otherwise of any trademark.
- The appearance of any proprietary name or reference in this document should not in itself be taken to imply a preference for one product over another unless specifically stated otherwise.
- The amounts and specifications shown in this presentation are a rough estimation and might not exactly match the ones that will be published for the tender


# Purpose of the GP Servers workstream

*”To provide all the server infrastructure stack from the bare metal up to the Application Layer, that is, to provide systems, either virtual or physical, ready to accept and run applications or services”*

# ECMWF will require the following kind of servers



Baremetal Servers



Hypervisors




Hyperconverged Nodes

Virtual Platform

ECMWF will require servers for:

Baremetal  
Servers



**ECPDS**

ECMWF Acquisition and  
Dissemination System

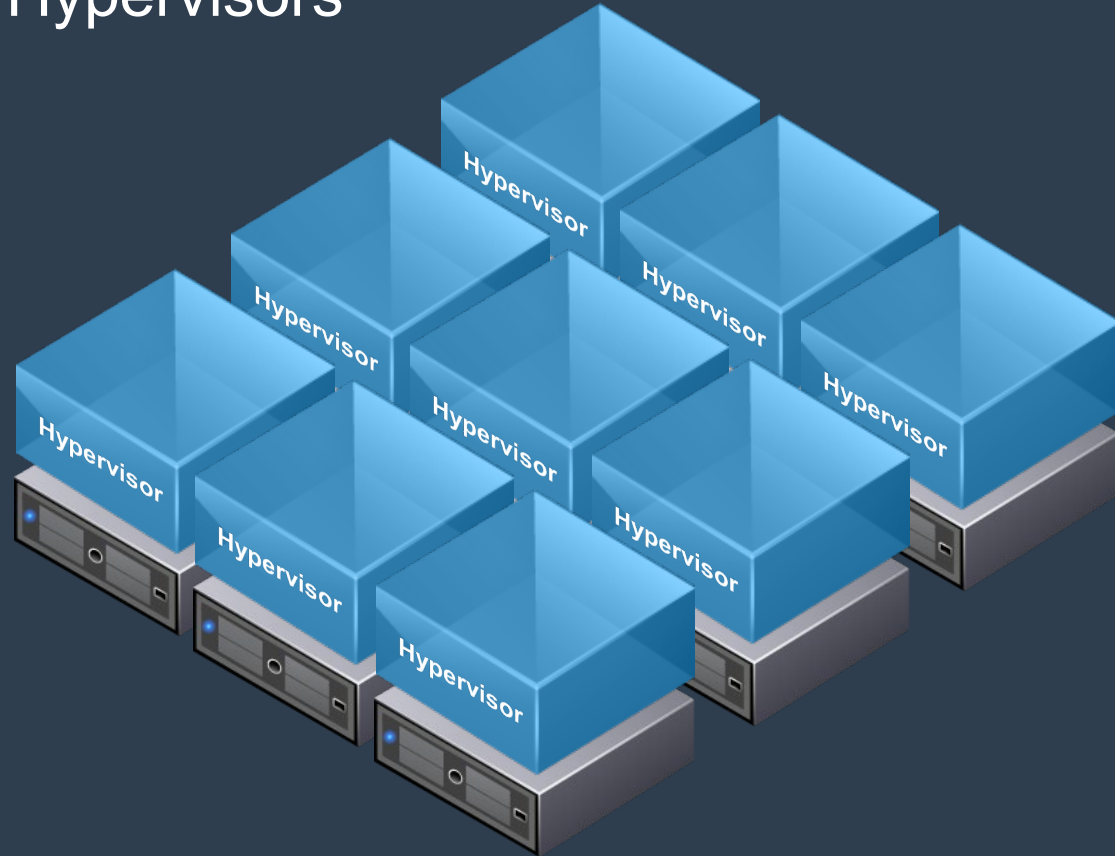
**DHS**

Data Handling System

Other ancillary systems

ECMWF will require servers for:


Hypervisors



**Virtualisation**  
General Purpose  
(Regular)  
Virtual Platform

ECMWF will require servers for

Hyperconverged  
Infrastructure



**Virtualisation**  
General Purpose  
(Regular)  
Virtual Platform

**VDI**  
Virtual Desktop  
Infrastructure

# Common requirements for all servers

- 19" standard rack compatible
- Intel-compatible CPU  
(Aiming at Silver-equivalent) and upper classes
- Networking: 25G default although 10G might be used
- Integrated Remote Management (ILO, iDrac, IMRC...) with prefailure warning
- Hotswappable & redundant fans and power supplies



# Common requirements for all servers

- Support and maintenance:  
3/3/3 (3-year parts, 3-year labour, 3-year onsite support)  
24/7 oncall, 4-hour response time (action)
- Integration might be required:  
OS preinstalled + management networking preconfigured.

# Baremetal Servers

- ECPDS (ECMWF Production Data Store)

Multi-purpose repository articulated around three strategic data-related services:

- Data Dissemination: automatic distribution of meteorological products to our Member States and other end-users.
- Data Acquisition: automatic discovery and retrieval of observational data from Data Providers.
- Data Portal: pulling of meteorological products and pushing of observational data initiated by remote sites.

Requirements:

50+ high-end servers with large, hybrid internal storage (~50TB each), no HBAs

# Baremetal Servers

- **DHS (Data Handling System)**

In DHS users can store and retrieve data that is needed to perform and research weather modelling, and mining of weather data. It comprises two main systems:

- MARS: Meteorological Archival and Retrieval System. Abstraction engine that allows staff and applications to access the meteorological data that has been collected or generated at ECMWF for more than 30 years
- ECFS: Provides users with a logical view of the file systems. Used for data not suitable for storing in MARS

Supporting ECFS and most of MARS is an underlying file archiving component, High Performance Storage System (HPSS), in which data is kept and managed

Requirements:

- 200+ servers with different specs, ranging from small 1U mid-range servers to 4U high-end servers, most of them with FC or SAS HBAs. It mostly uses SAN storage.

# Baremetal Servers

## Other servers

There are other servers that belong to some subsystems that, in total, are less than 50. They provide services for acquisition and pre-processing, networks, and NAS storage controllers.

Requirements are either 1U servers with FC HBAs and 2U servers with large local storage, mainly.

# Virtualisation Platform

- Holds the rest of GP services

Job control, Remote access, Directory Services, Web Services (www-apps-api-eccharts), Monitoring, Email relays, CMS, Software development Document Management System and archiving, Systems Configuration and Automation, etc...

- vSphere is the technology currently used at ECMWF

- Comprises GP Regular, VDI and GP HCI

- GP Regular runs on normal hypervisor nodes
- VDI runs on Hyperconverged nodes
- GP HCI runs on Hyperconverged nodes

- Over 900 VMs (400 production – 200 dev/test – 300 VDI)

# Virtualisation Platform – GP Regular

We are looking at small 1u servers:

- Medium-high CPU, large RAM (1TB)
- Small local storage for hypervisor OS
- Making use of shared IP storage (NAS) – 2 x dual NICs
  - This NAS (NFS) storage is specified and provided by the Storage Workstream
- Roughly 40 servers

# Virtualisation Platform – VDI

- An extra virtualisation platform is needed for the Virtual Desktop Infrastructure.
- The purpose is to hold ~300 Linux and ~100 Windows Virtual Desktops
- GPU hardware acceleration is required (NVIDIA)
- Full flash storage – up to 50TB usable
- We are looking at hyperconverged nodes like vxRails, Simplivity or DIY VSAN Ready nodes.
- Initially the requirements are ~14 servers

(High spec CPU, 1TB RAM, NVIDIA GPU, 2TB FF internal storage)

# Virtualisation Platform – GP HCI

- Objective: to have a bootstrap virtual infrastructure that can be built in no time from “day 0” at Bologna Datacentre
- Hybrid storage, up to 100TB.
- We are looking also at any kind of Ready Nodes like:
  - DIY ESXi nodes with VSAN
  - Dell vXRails
  - HPE Simplivity

(Less than 20 nodes, hybrid storage, medium-high spec CPU, 1 TB RAM, up to 5TB local space.)



# Summary

		Amount	Specs
Baremetal	ECPDS	~ 50	- High end CPU, ~50TB local storage, no HBAs
	MARS	~ 200	- Mostly 1U mid-range servers - Also 2U/4U high-end servers, large local storage - All with FC or SAS HBAs
	REST	~ 50	- 1U servers with FC HBAs - 2U servers with large local storage - 50/50
Virtual Platform	GP Regular	~ 50	1U servers, medium-high spec CPU, 1TB RAM, small local storage (OS), no HBAs
	VDI	~ 14	2U servers, high spec CPU, 1TB RAM, large flash storage, GPU card, no HBAs
	GP HCI	~ 20	2U servers, medium-high CPU spec, 1TB RAM, large hybrid storage, no HBAs

Q&A  
Grazie mille!