All information contained herein are for discussion purposes only and shall not be considered a commitment on the part of CERN or the buyers group.

addestino
innovation management

Context & Objectives

The Innovation Platform brings together

- (public) **demand/buyer**'s side
- and the (private) **supplier**'s side

To formulate a detailed answer on following key questions:

- What is innovation potential from **end-user perspective**?
- What is innovation potential from **techno/market point-of-view**?
- What are the priorities & criteria for **potential prototype**?

Hence to give guidance about rest of PCP trajectory

Prioritization grid: value vs. risk

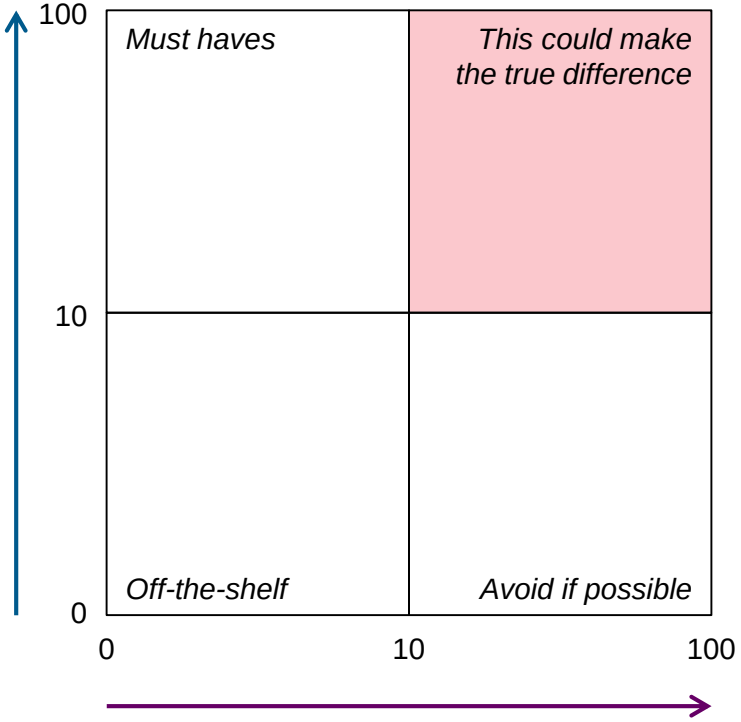

Value
(Innovation potential from end-user perspective)

Risk
(Innovation potential from technological/market point-of-view)

Innovative for end-user
+
Technological/market uncertainty

Approach

Done	1	Create list of **uses cases**
Yesterday	2	**Prioritize** these uses cases • I.e. innovation potential from end-user perspective
Yesterday	3	Build "**reference architecture**"
Today	4	Identify sources of technological **risk**
Today	5	**Prioritize** these risks • I.e. innovation potential from techno/market perspective
	6	Document findings and list priorities/criteria for **prototype**

addestino
innovation management

Structure of use cases

As a <Person>

I can <do something with the hybrid cloud>

So that <real value is created for me>

List of use cases (1/7)

1: Orchestration (13):
As IT Manager
I can provision resources (CPU, Storage, Network) through open APIs
So that transparently expanding my existing computing infra into Cloud is possible

2: Monitoring and Dashboards (13):
As IT Manager
I can extract in real-time (< 10 min) resource usage metrics,
through a dashboard or via APIs
So that I can see the allocated resources and adjust these if necessary

3: Reporting and accounting (13):
As IT Manager or Procurement Service
I can have a breakdown of the used resources over a period of time
So that an assessment of the accumulated cost of the resources consumed is possible

4: Service level agreements (13):
As IT Manager
I can have a set of standard service definitions (Scope, Responsibilities, KPIs, etc.)
So that service level is defined, understood, agreed between parties for each service

List of use cases (2/7)

5: Network IP addresses (13):

As IT Manager

I can use Public IPv4 addresses (or equivalent addressing techniques) from within the private cloud

So that majority of applications such as databases and remote read-only file systems (e.g. CERN CVMFS) that do no yet support IPv6, can be used.

6: Ordering, billing and invoicing mechanisms (13):

As IT Manager or Procurement Service

I can pay for and can get a receipt of the paid services

So that total expenditure of the used services is documented

7: Performant WAN connectivity via GEANT (13):

As IT manager

I can efficiently move and access data from local resources to commercial clouds which are connected to GEANT at a speed of 40 to 100 Gbits/s

So that I have a seamless, performant and reliable computing experience

8: Cloud economic models and procurement procedures of public organizations (20):

As Procurer

I can choose from different commercial delivery models e.g. reservation-based, spot-market, night-time bulk processing…

So that I can experiment with these and select the best model for my project

List of use cases (3/7)

9: Data privacy (13):
As IT Manager
I can restrict access to confidential data to specific users (e.g. LHC experiments)
So that data access can be restricted to an identified user community

10: Pay-as-you-go I (13):
As IT manager and Procurement Officer
I can have credits per cloud provider
so that I can see how many credits have been consumed for each cloud provider

11: Pay-as-you-go II (13):
As IT manager and Procurement Officer
I can refill the credits in my account and continue using cloud provider resources
so that I don‘t need to undertake an new procurement procedure and run the risk that contract conditions have changed

List of use cases (4/7)

12: Combination of firms/Sub-contracting (13):
As IT manager and Procurement Officer
I can expect suppliers that depend on subcontractors to have agreements with each subcontractor to provide at least 30% of the total required capacity
so that uninterrupted service can be ensured

13: Federated AAI and Credential Translation schemes (for resource mngmnt) (20):
As IT Manager or user
I can use my own credentials
So that the IaaS can easily be accessed and managed

14: Helpdesk and Computer Security Response Teams (13):
As IT Manager
I can use a common Service Desk and Security Response Team
So that reporting and troubleshooting operational/security problems across the full infrastructure can be guaranteed

15: Alerts (13):
As IT Manager
I can receive meaningful notifications and alerts about the resources I'm using
So that I have prompt information about infrastructural problems

List of use cases (5/7)

16: Efficient data access (40):
As a user
I can rely on the cloud provider to implement optimal strategies for handling very large datasets (e.g. via persistent caching, pre-staging)
So that I can reach high performance at an optimal cost (storage, network, CPU…)

17: Seamless service (20):
As end-user (scientist)
I can use the service provided by multiple providers seamlessly
So that I don't to be aware that the service is being provided by multiple providers

18: Metrics for aggregate performance, benchmarking (20):
As IT Manager
I can retrieve performance and benchmark metrics
So that effective end-to-end delivered service capacity & performance can be assessed

19: HPC (some 13, some 20):
As Procurer or IT Manager
I can in addition to typical configurations, also provision
large memory, high core counts and/or larger/more performant local storage via high performance backpanel
So that some end users are be able to run their massively parallel applications

List of use cases (6/7)

20: Multi-cloud management frameworks (20):
As IT Manager
I can use a single multi-cloud management framework via open APIs
So that the integration costs of managing multiple clouds can be reduced

21: Consultancy & training (8):
As a end-user or middleware developer
I can get help in migrating my application to the hybrid cloud
and I can get training on how to use this hybrid cloud
So that I can effectively use the platform for my research

22: Support (20):
As IT-Manager or system administrator
I can access adequately skilled support groups at the cloud provider
So that I can solve or act as coordinator/supervisor for the solution of complex problems running complex scientific workflows/applications

23: Container support (20):
As a user
I can develop, maintain, manage and run my containerized applications (with all local dependencies) in the hybrid cloud via standardized tools
So that I minimize the deployment effort for me and for the scientific workgroups

List of use cases (7/7)

24: Scaling of containers (20):

As a IT Manager & Middleware Developer

I can deploy massive amounts of containerized applications on clouds from various providers meeting requirements on stability & performance

So that I can easily and safely achieve massive scaling of scientific workloads with identical user and management experience

25: Transparent data access (40):

As a user

I can have an application running in the cloud which sees the full datasets/environment in the exactly the same way as on-site

So that I don't need to modify my application

26: Configure once, deploy everywhere (40):

As a user

I can write the deployment configuration (including possible performance strategies) of my application once

So that I can easily deploy my application to all cloud providers using the same configuration

List of technical requirements

- Core count per VM:
  - Most use cases require 1…8 cores
  - Some require 16 cores per VM (OCEAN, Euro BioImaging, PanCancer
  - Special requirement for massively parallel applications (128 cores or more) – high core count or fast interconnects (Infiniband?)
- RAM per core: most use cases fine with 2…4 GB
- Local storage at VM: Some requirements for large capacity and high IOPS
- Images
  - Linux (CentOS 6+, SL6/7, Debian)
- VM life times:
  - Some fraction of VMs expected to be stable over whole project phase,
  - Others could be short-lived (e.g. for 12 hours every day);
  - Life times of much less than 12 hours probably not very useful
- Workload is bound by budget not demand
- Minimum total capacity needed during implementation & sharing phases:
  - Prototype: 10’000 cores & 1 PB during approx. 6 months in 2017
  - Pilot: 20’000 cores, 2 PB during approx. 12 months in 2017-2018
- Shared storage as block or clustered or NAS storage

Planning Poker - Value

Planning poker is a 'best practice' methodology to estimate e.g. added value, level of complexity, required implementation effort, risk. The technique is based on domain expert evaluation and attaining consensus.

The cards:

The interpretation when assessing value:

0-2	**Could not care less**, not needed.
3-5	**OK, why not**, could perhaps bring some value
13	**Must have.** Solution should have this requirement
20-40	**WOW**, home run, this would really shift the boundaries !
?	**Don't know**, no experience with this subject..

addestino
innovation management

Planning Poker - Risk

Planning poker is a 'best practice' methodology to estimate e.g. added value, level of complexity, required implementation effort, risk. The technique is based on domain expert evaluation and attaining consensus.

The cards:

The interpretation when assessing risk:

0-2	**No worries**, of-the-shelf standard solutions exist.
3-5	**A frequent problem**, certainly solvable.
13	**Absolutely not a standard problem.** A solution requires important choices, thorough elaboration, and specific expert effort. Success can be achieved with significant time and effort.
100	**Impossible**, requires breakthroughs of physical laws.
?	**Don't know**, no experience with this subject..