A Novel Approach for Performance Characterization of IaaS Clouds

Sándor Ács, Márk Gergely, Zsolt Németh
acs.sandor@sztaki.mta.hu

Outline

- IaaS Clouds
- Performance and benchmark questions
- Performance characterization with hierarchial fuzzy inference system
- Case study
- Results and future works

IaaS Clouds

- A technology that helps to manage the big computational and storage capacity
- Public clouds promise flexibility, scalability, (SLA based) high availability and pay-as-you-go option for its users
- A cloud infrastructure provides efficient resource-management, lower operational costs for its maintainers

Classification

Service models:

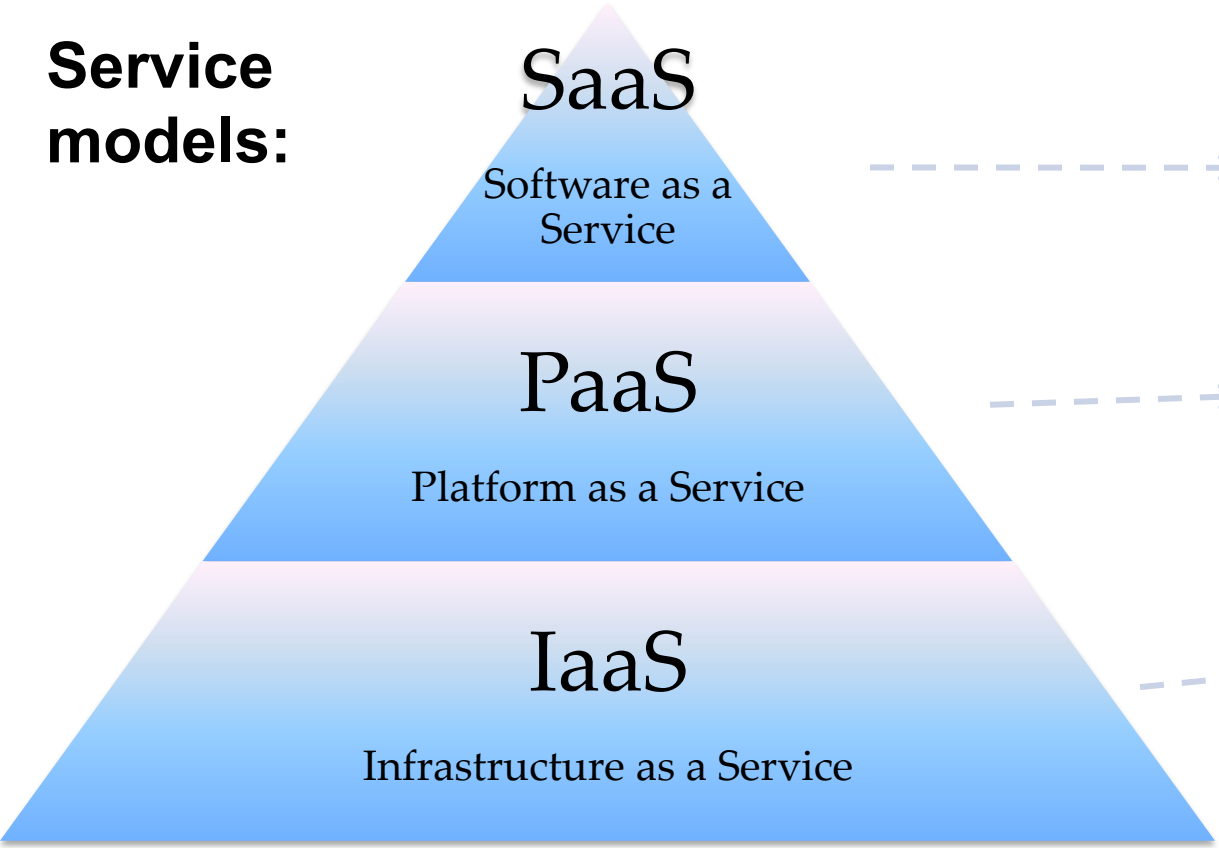

For Example: twitter, flickr …

Google App Engine

Amazon EC2, SZTAKI Cloud …

Deployment models: Private, public and hibrid

Driving forces

- Provider's needs
  - Improve resource utilisation
  - Use available resources efficiently
  - Save energy
  - Decrease cost
- Consumer needs
  - Cost effectiveness
  - Easy access to resources
  - Pay-as-you-go
  - No initial investments
  - Self-service

Iaas Clouds

- An economic model for resource
  - Provisioning
  - Management
- Service
  - Realizes utility computing
  - Elastic
  - Increased availability and reliability
  - Improved accessibility – ease of use
  - Reduced cost
  - Reduced energy consumption
- Must be comparable
  - To each other
  - Price vs value

What is the value of a service?

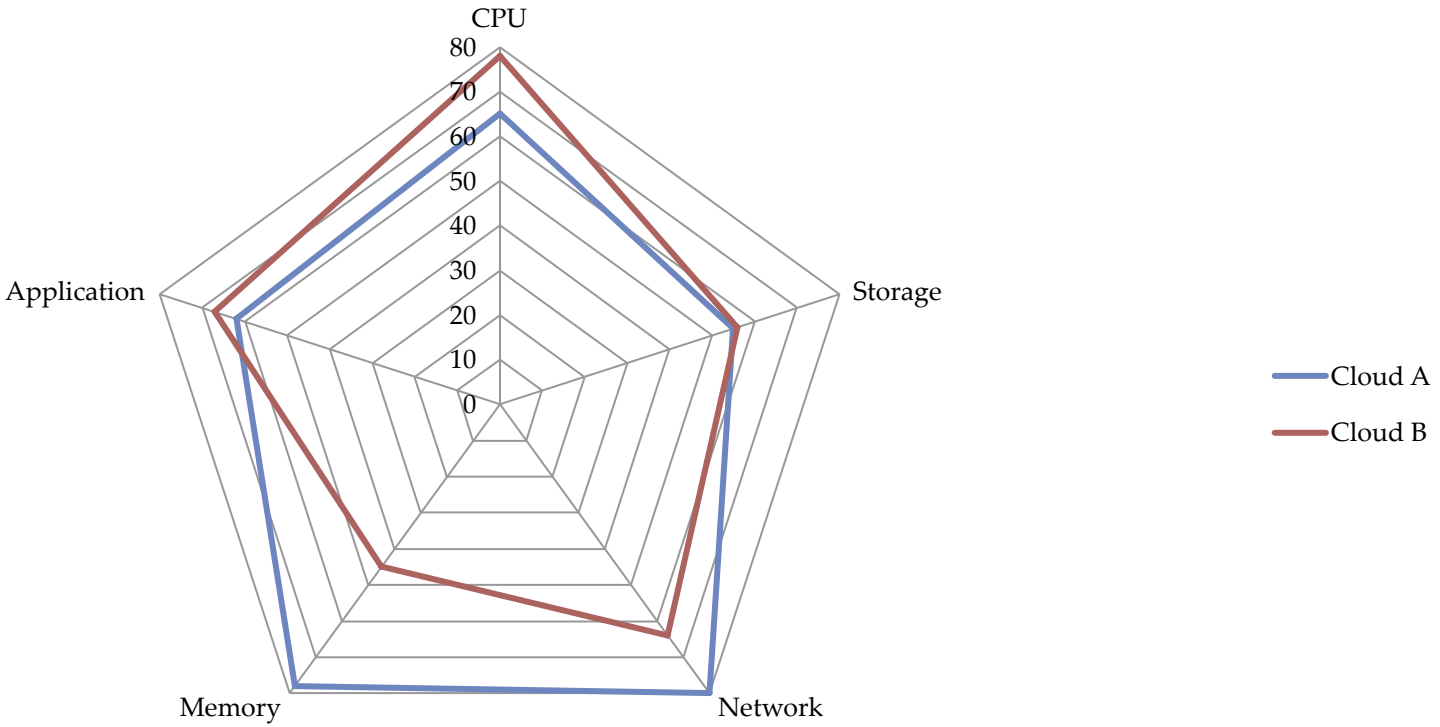

What is the value of an aspect?

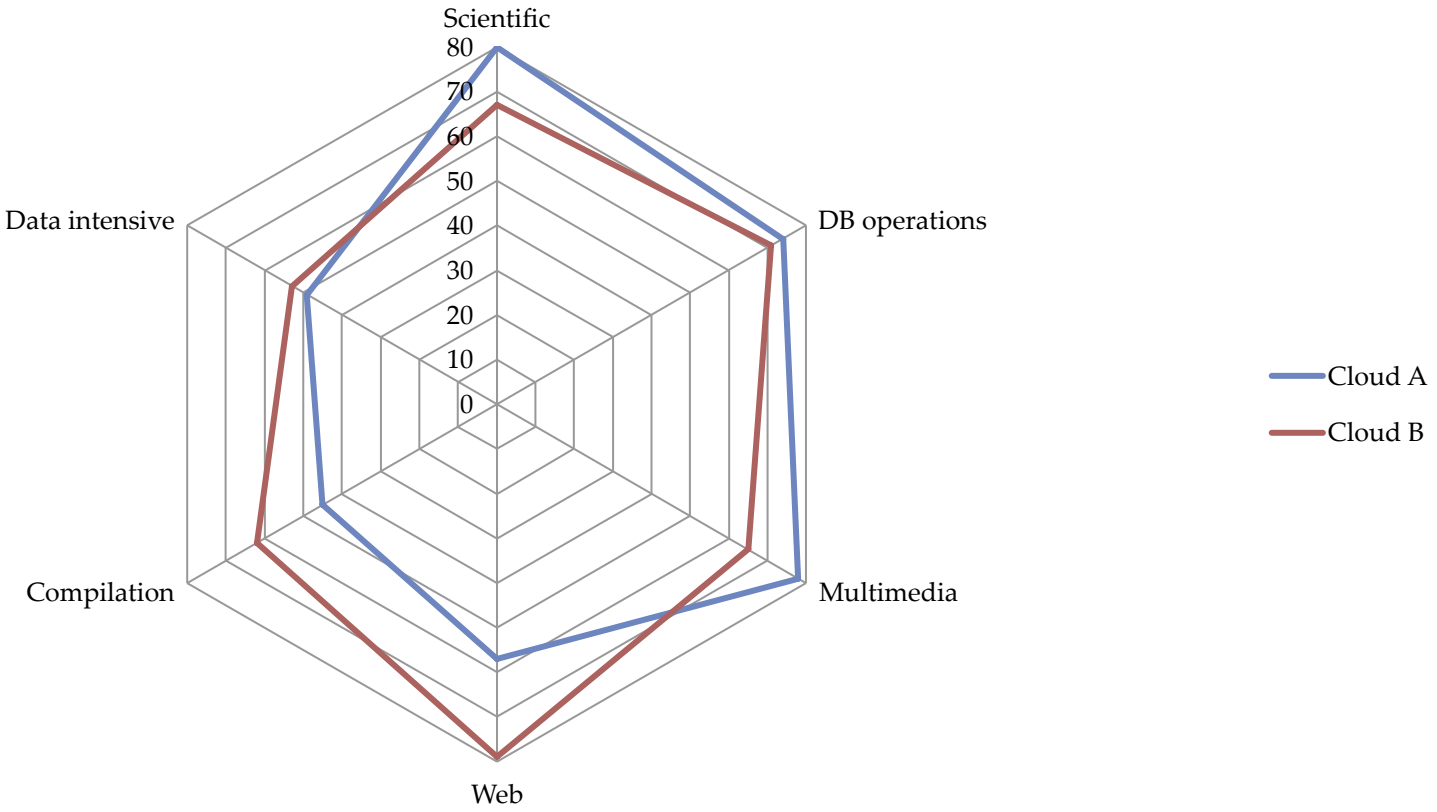

Benchmarking

Benchmarks

Since 2009, we h

Select the

- Performa
  - Aggre
  - CPU
  - Multi
  - IO Be
    - Di
    - M

4 service credits remaining | Buy Token | Apply Token

Service	Server	Data Center	Test Date	Test Score	Test Date	Test Score	Test Date	Test Score	Test Date	Test Score	Geekbench Score
Amazon EC2	ec2-us-east.linux.cc1.4xlarge-raid0x4-ebs	VA, US	03/06/2011	13830							13830
Amazon EC2	ec2-us-east.linux.cc1.4xlarge-raid0-local	VA, US	03/06/2011	13644							13644
Amazon EC2	ec2-us-east.linux.cg1.4xlarge	VA, US	02/28/2011	13434							13434
Amazon EC2	ec2-us-east.linux.cc1.4xlarge-local	VA, US	03/06/2011	13329							13
Amazon EC2	ec2-us-east.linux.cc1.4xlarge	VA, US	09/07/2010	12306	02/28/2011	13543					129
Amazon EC2	ec2-us-east.linux.m2.4xlarge	VA, US	04/19/2010	5877	02/26/2011	3537					
Amazon EC2	ec2-us-east.linux.m2.2xlarge	VA, US	04/19/2010	5163	03/20/2011	3534					4348.5
Amazon EC2	ec2-us-east.linux.c1.xlarge	VA, US	04/11/2010	5118	03/19/2011	2609					3863.5
Amazon EC2	ec2-us-east.linux.m2.xlarge	VA, US	04/19/2010	4049	09/07/2010	3952	03/19/2011	3377	05/25/2011	4057	3858.75
Amazon EC2	ec2-us-east.linux.m1.xlarge	VA, US	04/19/2010	4256	03/19/2011	2835					3545.5
Amazon EC2	ec2-us-east.linux.m1.large	VA, US	04/19/2010	3092	03/21/2011	2596					2844
Amazon EC2	ec2-us-east.linux.t1.micro	VA, US	09/09/2010	2568	03/21/2011	2802					2685

feedback feedback

Problem analysis

- Metrics
  - Not standardized
  - Very large dimension
  - Not comparable
- Virtualization
  - Split or merged physical resources to accommodate virtual machines
  - Different instance types
  - Multi-tenancy
- Benchmarking
  - Not applicable for comparison
  - Physical infrastructure vs. virtual machine instances
- There are no well established performance analysis techniques for IaaS clouds
- Performance analysis, performance evaluation of a cloud
  - Not applicable here rather ‘characterization’ of VM instances
  - A relative number for capturing performance

Concept

- Characterize instances
  - Not “benchmarking the cloud”
  - But benchmarks are used in the procedure
- Qualitative characterization
  - Uniform information
- Aggregate metrics
  - Decrease dimensions
- Provide an automated mechanism to perform characterization

Performance characterization with hierarchial fuzzy inference system

Qualitative metrics

- Metrics cannot be aggregated
  - They are not additive
  - Different dimensions
  - Not comparable
  - Not related

CPU A, frequency f_1, memory M_1, bandwidth b_1

vs

CPU B, frequency f_2, memory M_2, bandwidth b_2

- Quantitative metrics are transformed into qualitative
  - Operations on symbolic values

CPU A, “high” frequency, “large” memory, “low” bandwidth

vs

CPU B, “medium” frequency, “medium” memory, “medium” bandwidth

Fuzzy characterization

- Fuzzy values
  - transform quantitative information into qualitative
  - concise, readable, interpretable and comparable in an easy way
- Fuzzy inference
  - Simple readable statements transform fuzzy sets
  - If A is „good" and B is „excellent" then C is "good"
  - Rules and knowledge base represent expert opinion
- De-fuzzification
  - The result is a fuzzy set
  - Transformed into a comparable value by CoG

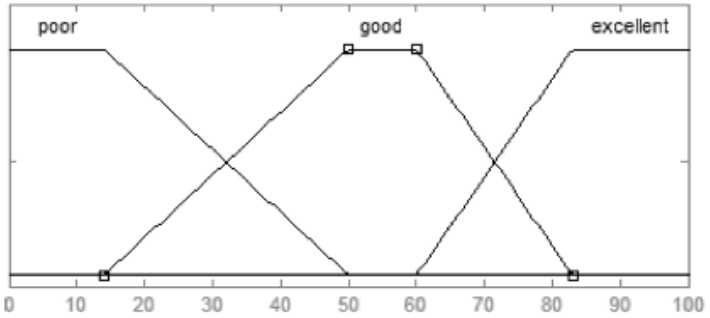

Architecture to gather raw data

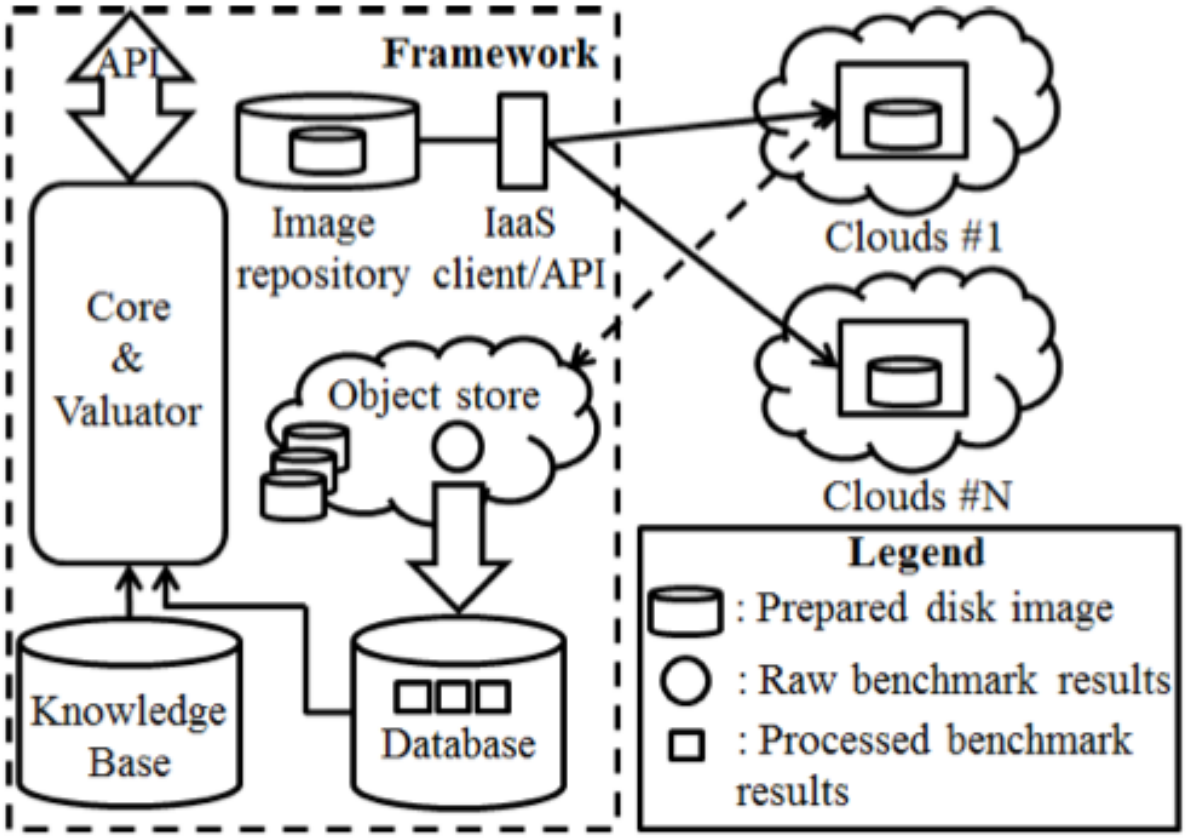

Raw data

- CPU performance
  - 11 tests
- Storage I/O
  - 6 tests
- Memory I/O
  - 3 tests
- Application
  - Compilation: 2 tests
  - Compression: 4 tests
  - Encoding: 5 tests
  - Database
  - Web: 2 tests

Fuzzy characterizaton

- If raw data is processed by fuzzy inference
  - Very high dimension
  - Rule become very complex
- CPU, memory, disk and network
  - captured by 157 parameters
  - number of corresponding fuzzy rules n^{157}
    - where n is the number of fuzzy sets (~granularity of rules)

Hierarchial Fuzzy Inference System (HFIS)

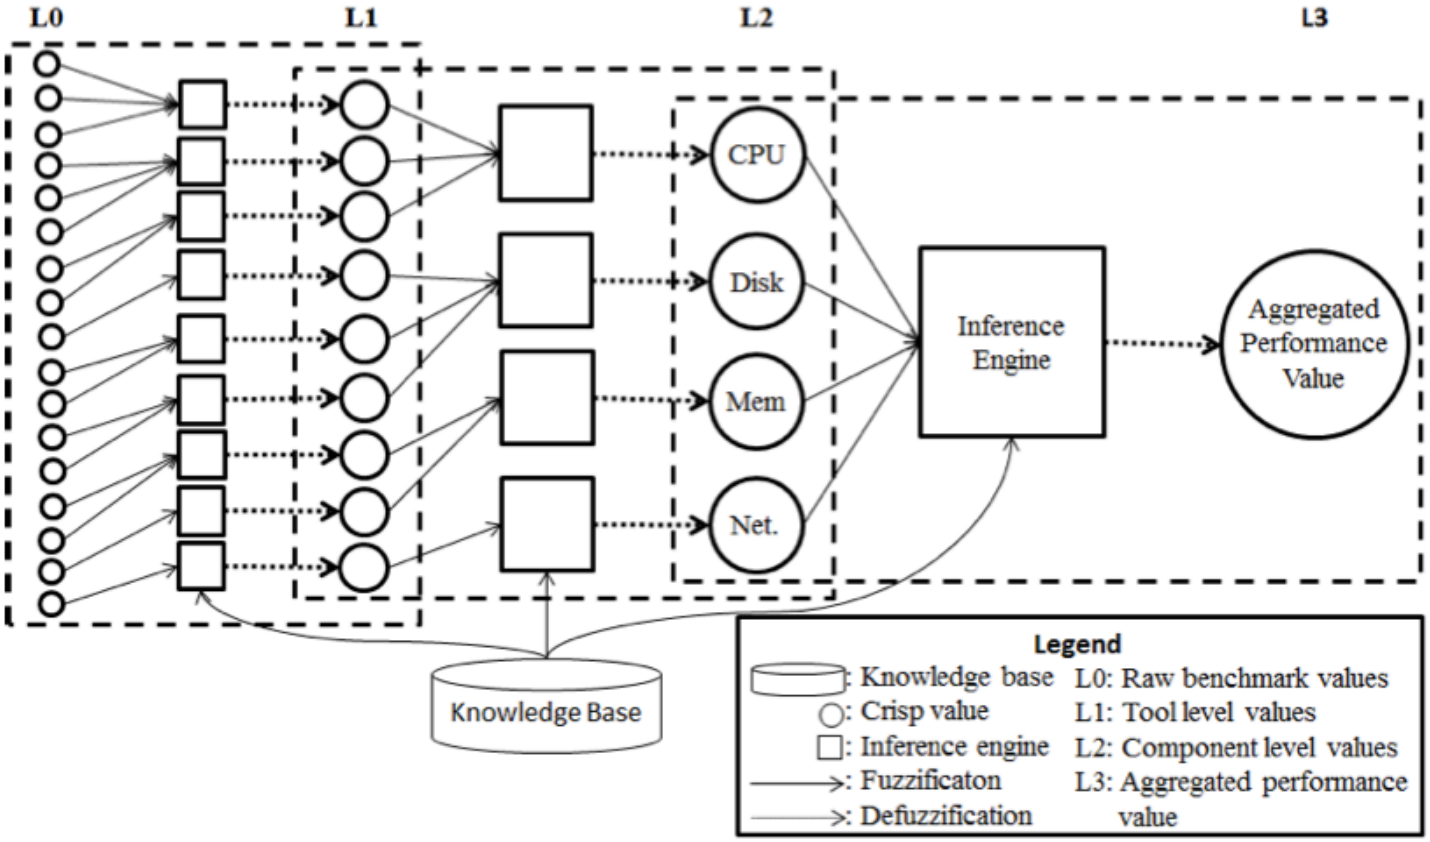

Hierarchial Fuzzy Inference System (HFIS)

- HFIS advantage
  - the number of rules is greatly reduced
  - the number of rules increases only linearly with the number of input variables
- HFIS divides the inference into stages
  - a subset of input variables produce intermediate results
  - these results are taken as inputs in subsequent stages
  - the number of variables in an inference stage is controlled
  - the intermediate results may also possess interpretable meaning
- Categorized the parameters according to the four main aspects
  - established sub-categories within each
  - input parameters to an inference stage do not exceed 7
  - overall number of rules in the system is bounded by $c * n^7$
    - c is the number of inference stages

Case study

- Can we compare Amazon EC2 and SZTAKI Cloud?
  - Disk I/O subsystem

Fs_mark	1000 Files, 1MB size	5000 Files, 1MB Size, 4 Threads	4000 Files, 32 Sub Dirs, 1MB Size	1000 Files, 1MB Size, No Sync / FSync
SZTAKI [Files/s]	58.55	93.87	68.33	132.63
Amazon [Files/s]	38.87	49.73	40.83	119.07

Dbench	1 client	6 clients	12 clients	48 clients	128 clients	256 clients
SZTAKI [MB/s]	113.385	225.64	242.08	220.02	185.53	134.38
Amazon [MB/s]	80.14	174.37	166.09	176.74	177.51	119.48

Case study

- Aggregated value by fuzzy inference

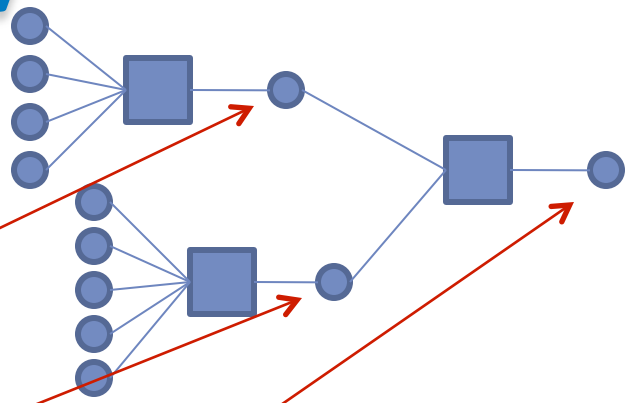

	DBench	fs_nark	Calculated value
SZTAKI [0-100]	72.16	62.55	62.98
Amazon [0-100]	62.06	43.6	47.05

- The two services are comparable as 62.98 : 47.05

Results and future works

- Feasibility study: comparing the Amazon EC2 and SZTAKI cloud standard instances
  - No, SZTAKI Cloud is not better than Amazon EC2
  - Rule base and knowledge base is complicated
- In the future, we plan to
  - Build a coherent and sound knowledge base
  - Investigate automatic mechanism to extract knowledge base
    - Refine qualities and relations (“what is excellent”, “what is good”, etc)
    - Tune the knowledge base a technology progresses

Thank you for the attention!

Questions?

Hierarchial Fuzzy Inference System (HFIS)

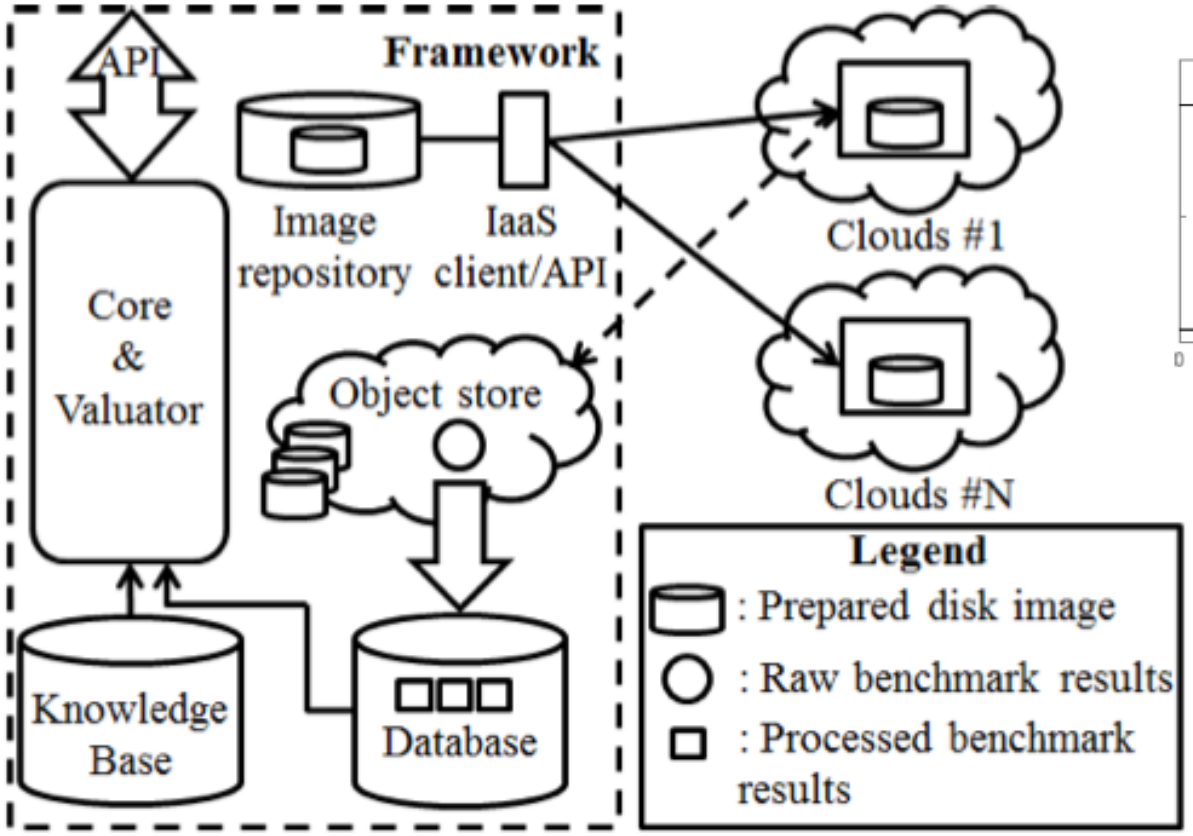

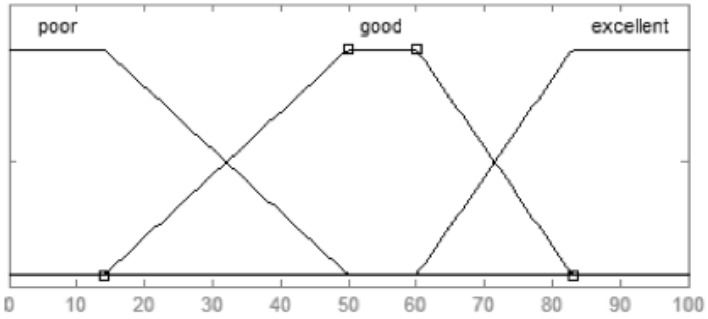

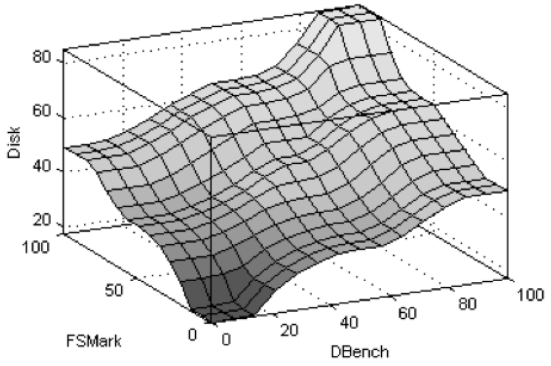

Current solution (2)

- **C-Meter**
  - Simple framework
  - Synthetic workloads via EC2 interface
- **Problem**
  - Only synthetic workloads
  - It uses an abandoned framework as a basis (GrenchMark)
  - The C-Meter seems to be abandoned as well