

## Cloud computing

- La definizione classica di riferimento è quella del National Institute of Standards and Technology (NIST) USA (<http://goo.gl/eBGBk>)

- In sintesi il Cloud computing si occupa di:

**Fornitura di tecnologia di informazione e comunicazione (ICT) come servizio**

David Salomoni Introduzione al Cloud Computing - 9 Dec 2013

11

David Salomoni

12

Introduzione al Cloud Computing - 9 Dec 2013

## Caratteristiche del Cloud

- **Self-service, on-demand**
  - Il cliente chiede autonomamente ciò che gli serve, quando gli serve (e sperabilmente lo ottiene).
- **Accesso attraverso la rete**
  - Assume che una rete (Internet o intranet) sia disponibile, normalmente a banda larga.
- **Pool di risorse**
  - L'utente non si preoccupa di conoscere i dettagli delle risorse, che sono gestiti dai Cloud resource provider.

### Elasticità

- L'utente non si preoccupa di conoscere i dettagli delle risorse, che sono gestiti dalle necessità del cliente.

### Pagamento a consumo

- Il cliente paga solo per ciò che usa.

## Il focus sul “Service”

- Abbiamo visto che nella definizione di Cloud computing (“Fornitura di tecnologia di informazione e comunicazione come servizio”) il **servizio** nei confronti del cliente è parte essenziale.
- Il Cloud computing si può modellare infatti intorno a servizi legati principalmente a
  - Infrastruttura (**IaaS** → Infrastructure as a Service)
  - Piattaforma (**PaaS** → Platform as a Service)
  - Software (**SaaS** → Software as a Service)

## Una analogia: l'autonoleggio



- Self-service, on-demand
  - Prenotazione telefonica oppure online
- Rete
  - Estesa rete di autonoleggi in tutto il mondo
- Pool di risorse
  - Pensa l'autonoleggio a gestire sapere quante macchine gli servono
- Elasticità
  - Il numero di auto disponibili normalmente varia a seconda della richiesta
- Pagamento a consumo
  - Il cliente paga per il tempo in cui usa l'auto (e non pensa ad assicurazione, gomme, etc.)

Fonte: <http://goo.gl/cEa8M>

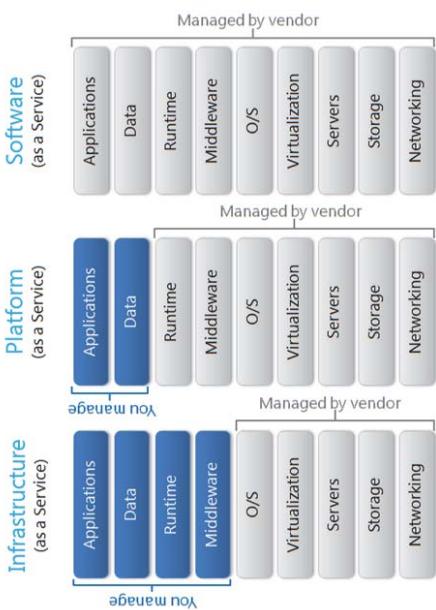
David Salomoni Introduzione al Cloud Computing - 9 Dec 2013

13

Introduzione al Cloud Computing - 9 Dec 2013

17

## Chi fa cosa?



Fonte: <http://go0o.gl/1imkR>

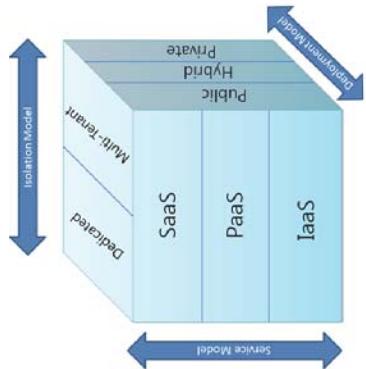
David Salomoni Introduzione al Cloud Computing - 9 Dec 2013 18

Fonente: <http://go0o.gl/1imkR>

David Salomoni Introduzione al Cloud Computing - 9 Dec 2013 19

## Aggiungiamo dimensioni

- Oltre i modelli di servizio, parti importanti per definire e capire il Cloud computing sono i modelli di:
  - deployment** (dove distribuisco i servizi)
  - isolamento** (come isolo i servizi)



Fonte: <http://go0o.gl/1imkR>

David Salomoni Introduzione al Cloud Computing - 9 Dec 2013 19

Fonente: <http://go0o.gl/1imkR>

David Salomoni Introduzione al Cloud Computing - 9 Dec 2013 19

## Deployment: i “tipi di Cloud”

- Cloud privata:**
  - L'infrastruttura viene fornita per un uso esclusivo da parte di una singola organizzazione. La gestione, l'operazione, la proprietà, la dislocazione della Cloud privata tuttavia può essere anche indipendente dall'organizzazione che la usa.
- Cloud di comunità (Community Cloud):**
  - L'infrastruttura è disponibile ad una comunità di organizzazioni che hanno uno scopo comune (ad esempio missione, requisiti di sicurezza, conformità a regole comuni, etc.)
- Cloud pubblica:**
  - L'infrastruttura è disponibile in generale al pubblico. La gestione può essere pubblica o privata. La dislocazione è presso il fornitore di servizi.
- Cloud ibrida:**
  - L'infrastruttura è una combinazione di due o più infrastrutture Cloud (private, di comunità o pubbliche) che sono collegate in modo da garantire forme di portabilità ad esempio di dati o applicazioni.

Fonte: <http://go0o.gl/1imkR>

David Salomoni Introduzione al Cloud Computing - 9 Dec 2013 20

David Salomoni Introduzione al Cloud Computing - 9 Dec 2013 21

## Isolamento

- I modelli di **isolamento** nel Cloud (spesso ignorati) sono importanti e si dividono in:
  - Infrastrutture dedicate
  - Infrastrutture “multi-tenant” (con diversi [tipi di] clienti)
- Il tipo di isolamento è importante per molti aspetti, come:
  - Segmentazione delle risorse
  - Protezione dei dati
  - Sicurezza delle applicazioni
  - Auditing
  - Disaster recovery

Fonte: <http://go0o.gl/1imkR>

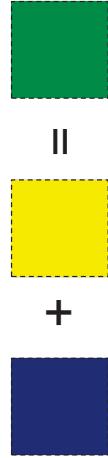
David Salomoni Introduzione al Cloud Computing - 9 Dec 2013 21

## Che differenza c'è...

- ... tra virtualizzazione e Cloud computing?



- Risposta:



David Salomoni

Introduzione al Cloud Computing - 9 Dec 2013

28

29

## Virtualizzazione?

- Il Cloud computing può anche essere fornito senza l'utilizzo di tecnologie di virtualizzazione.
  - Spesso tuttavia l'utilizzo di tecnologie di virtualizzazione consente di ridurre i costi operativi e in conto capitale.
  - Essere in grado di fornire molto rapidamente delle macchine virtuali non è comunque efficiente, se servono diversi mesi per effettuare il provisioning e l'installazione degli host fisici.
    - Inoltre, il tempo impiegato per la fornitura dello strato di virtualizzazione è recuperato dai risparmi associati al non dover utilizzare server fisici?
    - Importanza di avere tool di installazione, monitoring e accounting il più possibile automatizzati.
  - Ma che cosa si intende con virtualizzazione?**

David Salomoni

Introduzione al Cloud Computing - 9 Dec 2013

29

## Riassumendo: virtualizzazione vs. Cloud computing

### Cloud computing

- Installazione/reinstallazione di server o di applicazioni su VM di per sé **non è Cloud computing**.
- Verificare con le **5 caratteristiche del Cloud** mostrate precedentemente:
  - Self-service, on-demand → **NO** (tipicamente è un dipartimento IT che fornisce le VM)
  - Accesso attraverso la rete → **NO** (deployment limitato a "internal customers")
  - Pool di risorse → **Sì**
  - Elasticità → **NO** (tipicamente è un dipartimento IT che deve installare sistema operativo + software, e non necessariamente in modo scalabile)
  - Pagamento a consumo → **NO** (spesso il billing non viene fatto a consumo ma in modo tradizionale)
- Un esempio di virtualizzazione che non è Cloud?** (ma che è complementare al Cloud)

- Cloud Computing  
Part 1: What is Cloud Computing?

University of Padova  
Department of Mathematics

July 14, 2014

At a glance	Reference model	Economics	Open challenges
<h2>Summary</h2>			

At a glance	Reference model	Economics	Open challenges
<h2>Computing utilities - Vision</h2>			

- ① Cloud computing at a glance

### ② The NIST definition of cloud computing

- ③ Economics of the cloud

- ④ Open challenges

Cloud Computing - Part 1

2 / 20

At a glance	Reference model	Economics	Open challenges
<h2>Cloud - The term</h2>			

- One of the most diffuse views of cloud computing

*"I don't care where my servers are, who manages them, where my documents are stored, or where my applications are hosted. I just want them **always available** and access them from any device connected through Internet. And I am willing to pay for this service for **as a long as I need it.**"*

- strong similarities to the way we use other services, such as water and electricity
- made possible by the effective composition of several technologies, which have reached the appropriate maturity level:
- Web 2.0 technologies ⇒ Internet into a rich application and service delivery platform
- Service orientation ⇒ to deliver capabilities with familiar abstractions
- Virtualization ⇒ confers on cloud computing the necessary degree of customization, control, and flexibility for building production and enterprise systems.

Cloud Computing - Part 1

3 / 20

At a glance	Reference model	Economics	Open challenges
<h2>Computing utilities - Enablers</h2>			

- The term *cloud* has historically been used in the telecommunications industry:

- abstraction of e network in system diagrams
- became the symbol of the most popular computer network: the Internet

This meaning also applies to cloud computing, which refers to an Internet-centric way of computing. Internet plays a fundamental role.

*"Cloud computing refers to both the applications delivered as services over the Internet and the hardware and system software in the datacenters that provide those services."*

Touching on the entire stack:

- XaaS ⇒ everything as a service
- different component of a system:
  - delivered
  - measured
  - priced, as a service

Cloud Computing - Part 1

4 / 20

Cloud Computing - Part 1

5 / 20

At a glance	Reference model	Economics	Open challenges
<h2>Cloud Computing at a glance</h2>			

The utility-oriented nature of cloud computing:

*"A cloud is a type of parallel and distributed system consisting of a collection of interconnected and virtualized computers that are dynamically provisioned and presented as one or more unified computing resources based on service-level agreements established through negotiation between the service provider and consumers."*

Buyya et.al

Cloud Computing - Part 1

6 / 20

At a glance	Reference model	Economics	Open challenges
<h2>Cloud Computing - 5 essential characteristics</h2>			

- On-demand self-service
  - a consumer can unilaterally provision computing capabilities
  - doesn't require human interaction with each service provider
- Broad network access
  - capabilities available over the network
  - accessed through standard mechanisms
  - heterogeneous, thin or thick client platforms
- Resource pooling
  - provider's computing resources are pooled to serve multiple consumers
  - multi-tenant model
  - different physical and virtual resources dynamically assigned and reassigned according to consumer demand
  - customer generally has no control or knowledge over the exact location
  - may be able to specify location at a higher level of abstraction

At a glance	Reference model	Economics	Open challenges
<h2>Cloud Computing - NIST Reference Model</h2>			

*"Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."*

Composed of: 5 essential characteristics, 3 service models, 4 deployment models.



At a glance	Reference model	Economics	Open challenges
<h2>Cloud Computing - 5 essential characteristics</h2>			

7 / 20

Cloud Computing - Part 1

Cloud Computing - Part 1

9 / 20

At a glance	Reference model	Economics	Open challenges
Cloud Computing - 3 service models			

- Software as a Service (SaaS)
  - consumers use the provider's applications running on a cloud infrastructure
  - applications accessible from various client devices
  - possible exception of limited user-specific application configuration settings
- Platform as a Service ( PaaS )
  - consumers deploy onto the cloud infrastructure
  - consumers create or acquire applications using tools supported by the provider
  - consumers has control over the deployed applications and possibly configuration settings for the application-hosting environment
- Infrastructure as a Service (IaaS)
  - consumers provision computing resources (e.g., processing, storage, network)
  - consumers are able to deploy and run arbitrary software, including OS and apps
  - consumers have control over OS, storage, and deployed applications
  - possibly limited control of select networking components (e.g., host firewalls)

Consumers does **not** manage or control the underlying cloud infrastructure !!!

At a glance	Reference model	Economics	Open challenges
Cloud Computing - 4 deployment models			

...continued

- Public cloud
  - cloud infrastructure provisioned for open use by the general public
  - may be owned, managed, and operated by a business, academic, or government organization, or some combination of them
  - it exists on the premises of the cloud provider
- Hybrid cloud
  - cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public)
  - remain unique entities , but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds)

At a glance	Reference model	Economics	Open challenges
Cloud Computing - 4 deployment models			

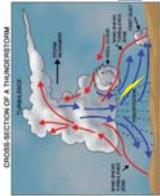
At a glance	Reference model	Economics	Open challenges
Cloud Computing - Part 1			

At a glance	Reference model	Economics	Open challenges
Cloud Computing - 4 deployment models			

At a glance	Reference model	Economics	Open challenges
Cloud Computing - Part 1			

## CS5412: ANATOMY OF A CLOUD

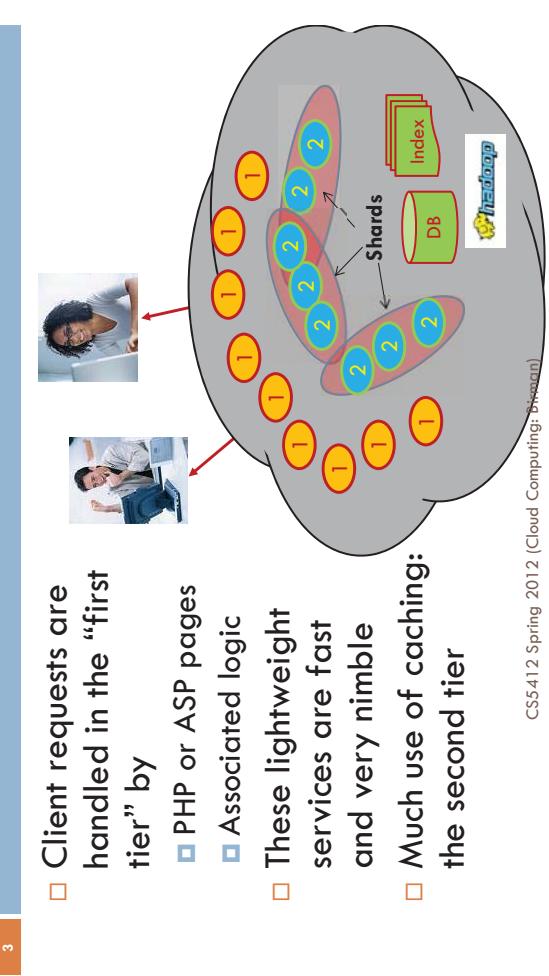
## How are cloud structured?



- 2 □ Clients talk to clouds using web browsers or the web services standards
  - But this only gets us to the outer “skin” of the cloud data center, not the interior
  - Consider Amazon: it can host entire company web sites (like Target.com or Netflix.com), data (AC3), servers (EC2) and even user-provided virtual machines!

CS5412 Spring 2012 (Cloud Computing: Birman)

## Big picture overview



CS5412 Spring 2012 (Cloud Computing: Birman)

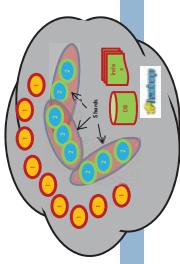
## Many styles of system

- 4 □ Near the edge of the cloud, focus is on vast numbers of clients and rapid response
  - Inside we find high volume services that operate in a pipelined manner, asynchronously
  - Deep inside the cloud we see a world of virtual computer clusters that are scheduled to share resources and on which applications like MapReduce (Hadoop) are very popular
- 5 □ We need to replicate
  - Processing: each client has what seems to be a private, dedicated server (for a little while)
  - Data: as much as possible, that server has copies of the data it needs to respond to client requests without any delay at all
  - Control information: the entire structure is managed in an agreed-upon way by a decentralized cloud management infrastructure

CS5412 Spring 2012 (Cloud Computing: Birman)

CS5412 Spring 2012 (Cloud Computing: Birman)

## What about the “shards”?



6

- The caching components running in tier two are central to the responsiveness of tier-one services
- The basic idea is to always use cached data if at all possible, so the inner services (here, a database and a search index stored in a set of files) are shielded from “online” load
- We need to replicate data within our cache to spread loads and provide fault-tolerance
- But not everything needs to be “fully” replicated. Hence we often use “shards” with just a few replicas

CS5412 Spring 2012 (Cloud Computing: Birman)

## Sharding used in many ways

7

- The second tier could be any of a number of caching services:
  - Memcached: a sharable in-memory key-value store
  - Other kinds of DHTs that use key-value APIs
  - Dynamo: A service created by Amazon as a scalable way to represent the shopping cart and similar data
  - BigTable: A very elaborate key-value store created by Google and used not just in tier-two but throughout their “GooglePlex” for sharing information
  - Notion of sharding is cross-cutting
  - Most of these systems replicate data to some degree

CS5412 Spring 2012 (Cloud Computing: Birman)

## Do we always need to shard data?

8

- Imagine a tier-one service running on 100k nodes
  - Can it ever make sense to replicate data on the entire set?
  - Yes, if some kinds of information might be so valuable that almost every external request touches it
  - Must think hard about patterns of data access and use
  - Some information needs to be heavily replicated to offer blindingly fast access on vast numbers of nodes
  - The principle is similar to the way Beehive operates.
    - Even if we don’t make a dynamic decision about the level of replication required, the principle is similar
    - We want the level of replication to match level of load and the degree to which the data is needed on the critical path

## And it isn't just about updates

9

- Should also be thinking about patterns that arise when doing reads (“queries”)
  - Some can just be performed by a single representative of a service
  - But others might need the parallelism of having several (or even a huge number) of machines do parts of the work concurrently
- The term sharding is used for data, but here we might talk about “parallel computation on a shard”

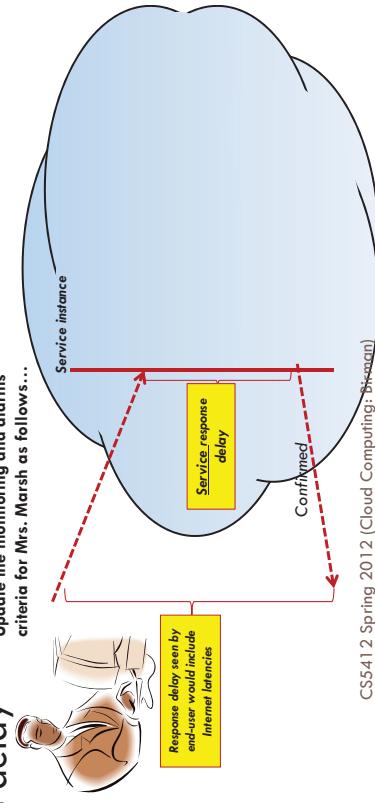
CS5412 Spring 2012 (Cloud Computing: Birman)

CS5412 Spring 2012 (Cloud Computing: Birman)

## What does “critical path” mean?

10

- Focus on delay until a client receives a reply
- Critical path is formed by actions that contribute to this delay



CS5412 Spring 2012 (Cloud Computing: Birman)

## What if a request triggers updates?

11

- If the updates are done “asynchronously” we might not experience much delay on the critical path
- Cloud systems often work this way
  - Avoids waiting for slow services to process the updates but may force the tier-one service to “guess” the outcome
  - For example, could optimistically apply update to value from a cache and just hope this was the right answer
- Many cloud systems use these sorts of “tricks” to speed up response time

CS5412 Spring 2012 (Cloud Computing: Birman)

## First-tier parallelism

12

- Parallelism is vital to speeding up first-tier services
- Key question:
  - Request has reached some service instance X
  - Will it be faster...
    - ... For X to just compute the response
    - ... Or for X to subdivide the work by asking subservices to do parts of the job?
- Glimpse of an answer
  - Werner Vogels, CTO at Amazon, commented in one talk that many Amazon pages have content from 50 or more parallel subservices that ran, in real-time, on your request!

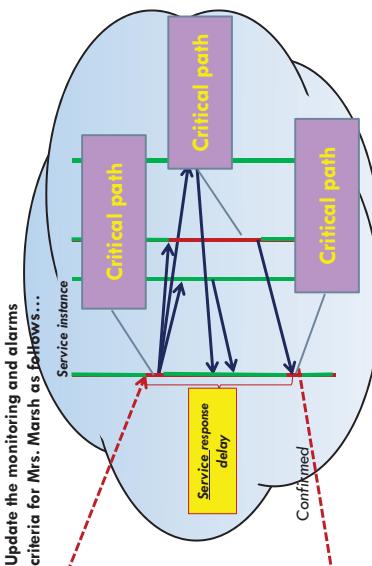
CS5412 Spring 2012 (Cloud Computing: Birman)

## What does “critical path” mean?

13

- In this example of a parallel read-only request, the critical path centers on the middle “subservice”
  - Update the monitoring and alarms criteria for Mrs. Marsh as follows...
    - Service instance
    - Critical path
    - Critical path
    - Critical path
  - Service response delay
  - Confirmed

Response delay seen by end-user would include Internet latencies



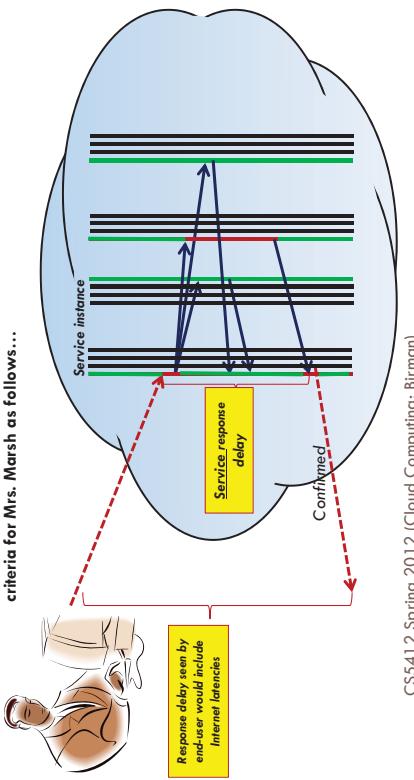
CS5412 Spring 2012 (Cloud Computing: Birman)

## With replicas we just load balance

But when we add updates....

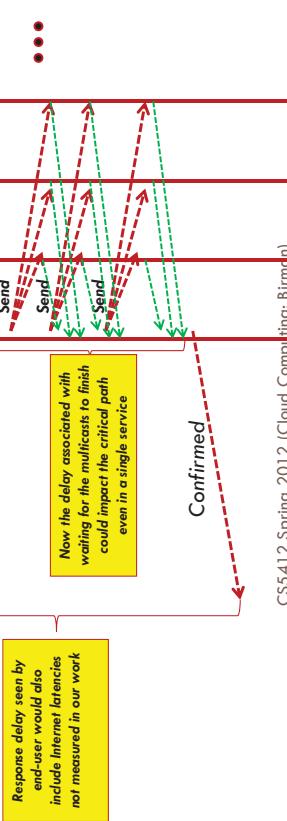
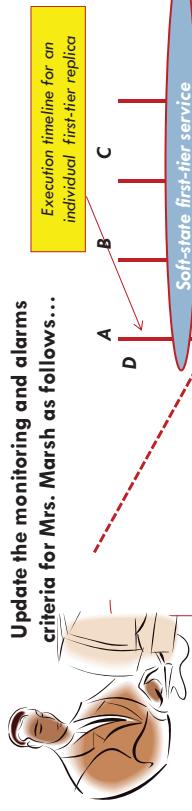
14

Update the monitoring and alarms criteria for Mrs. Marsh as follows...



CS5412 Spring 2012 (Cloud Computing: Birman)

15



CS5412 Spring 2012 (Cloud Computing: Birman)

Update the monitoring and alarms criteria for Mrs. Marsh as follows...



Response delay seen by end-user would also include internal latencies not measured in our work

Execution timeline for an individual first-tier replica

A  
B  
C  
D

Send

What about sending updates without waiting?

16

Several issues now arise

- Are all the replicas applying updates in the same order?
  - Might not matter unless the same data item is being changed
  - But then clearly we do need some "agreement" on order
  - What if the leader replies to the end user but then crashes and it turns out that the updates were lost in the network?
    - Data center networks are surprisingly lossy at times
    - Also, bursts of updates can queue up
  - Such issues result in inconsistency

CS5412 Spring 2012 (Cloud Computing: Birman)

## Eric Brewer's CAP theorem

17

- In a famous 2000 keynote talk at ACM PODC, Eric Brewer proposed that "you can have just two from Consistency, Availability and Partition Tolerance"
  - He argues that data centers need very snappy response, hence availability is paramount
  - And they should be responsive even if a transient fault makes it hard to reach some service. So they should use cached data to respond faster even if the cached entry can't be validated and might be stale!
  - Conclusion: weaken consistency for faster response

CS5412 Spring 2012 (Cloud Computing: Birman)

## CAP theorem

18

- A proof of CAP was later introduced by MIT's Seth Gilbert and Nancy Lynch
  - Suppose a data center service is active in two parts of the country with a wide-area Internet link between them
    - We temporarily cut the link ("partitioning" the network)
    - And present the service with conflicting requests
      - The replicas can't talk to each other so can't sense the conflict
      - If they respond at this point, inconsistency arises

CS5412 Spring 2012 (Cloud Computing: Birman)



CS5412 Spring 2012 (Cloud Computing: Birman) 20



## eBay's Five Commandments

21

- As described by Randy Shoup at LADIS 2008

*Thou shalt...*

1. Partition Everything
2. Use Asynchrony Everywhere
3. Automate Everything
4. Remember: Everything Fails
5. Embrace Inconsistency



## Is inconsistency a bad thing?

19

- How much consistency is really needed in the first tier of the cloud?
  - Think about YouTube videos. Would consistency be an issue here?
  - What about the Amazon "number of units available" counters. Will people notice if those are a bit off?
- Puzzle: can you come up with a general policy for knowing how much consistency a given thing needs?
- If

CS5412 Spring 2012 (Cloud Computing: Birman)

CS5412 Spring 2012 (Cloud Computing: Birman)

## Vogels at the Helm



22

- Werner Vogels, CTO at Amazon.com ...
- He was involved in building a new shopping cart service
  - The old one used strong consistency for replicated data
  - New version was build over a DHT, like Chord, and has weak consistency with eventual convergence



CS5412 Spring 2012 (Cloud Computing: Birman)

## James Hamilton's advice



23

- Key to scalability is decoupling, loosest possible synchronization
- Any synchronized mechanism is a risk
  - His approach: create a committee
    - Anyone who wants to deploy a highly consistent mechanism needs committee approval



CS5412 Spring 2012 (Cloud Computing: Birman)

## Consistency

24

- This weakens guarantees ... but

**Speed matters more than correctness**

CS5412 Spring 2012 (Cloud Computing: Birman)

## But inconsistency brings risks too!

25

- My rent check bounced?  
That can't be right!
- Inconsistency causes bugs
- Clients would never be able to trust servers... a free-for-all



**Consistency technologies just don't scale!**



- Weak or "best effort" consistency?

- Strong security guarantees demand consistency
- Would you trust a medical electronic-health records system or a bank that used "weak consistency" for better scalability?

CS5412 Spring 2012 (Cloud Computing: Birman)

CS5412 Spring 2012 (Cloud Computing: Birman)

## Puzzle: Is CAP valid in the cloud?

26

- Facts: data center networks don't normally experience partitioning failures
  - Wide-area links do fail
  - But most services are designed to do updates in a single place and mirror read-only data at others
  - So the CAP scenario used in the proof can't arise
- Brewer's argument about not waiting for a slow service to respond does make sense
  - Argues for using any single replica you can find
  - But does this preclude that replica being consistent?

CS5412 Spring 2012 (Cloud Computing: Birman)

## What does “consistency” mean?

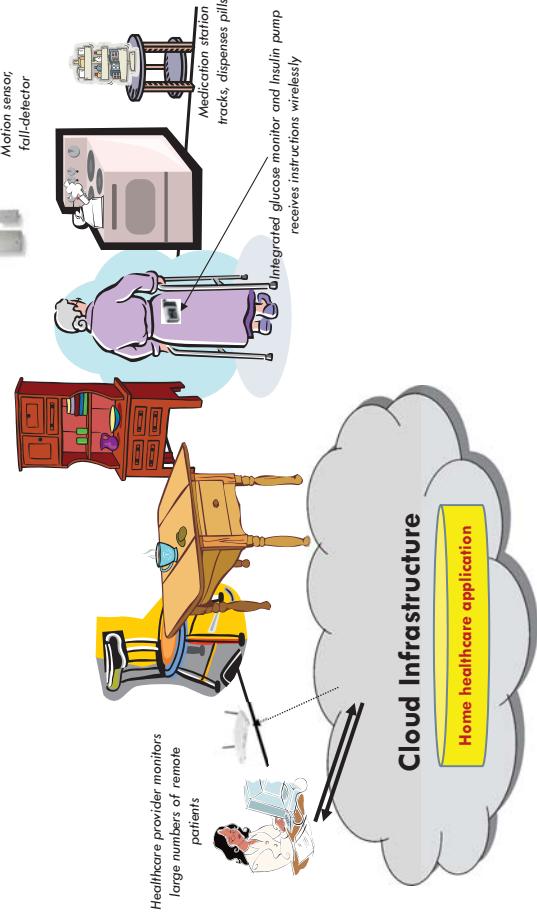
27

- We need to pin this basic issue down!
- As used in CAP, consistency is about two things
  - First, that updates to the same data item are applied in some agreed-upon order
  - Second, that once an update is acknowledged to an external user, it won't be forgotten
- Not all systems need both properties

CS5412 Spring 2012 (Cloud Computing: Birman)

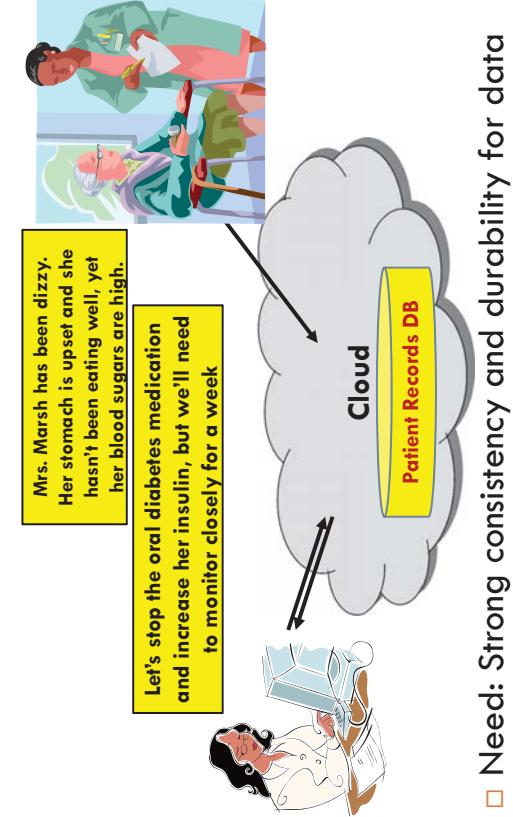
## What properties are needed in remote medical care systems?

28



## Which matters more: fast response, or durability of the data being updated?

29

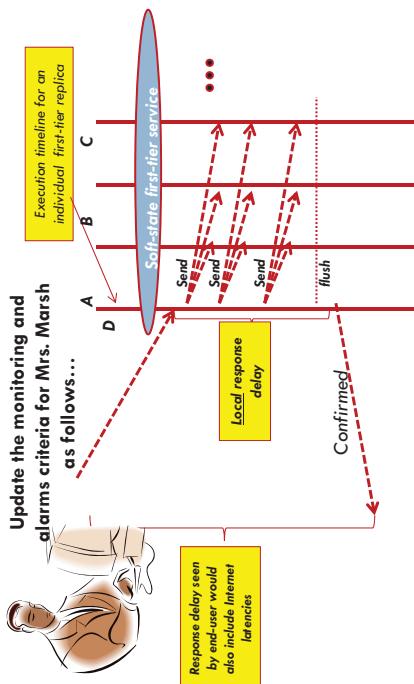


- Need: Strong consistency and durability for data

## What if we were doing online monitoring?

## Why does monitoring have weaker needs?

30



- An online monitoring system might focus on real-time response and value consistency, yet be less concerned with durability

31

- When a monitoring system goes “offline” the device turns a red light or something on
  - Later, on recovery, the monitoring policy may have changed and a node would need to reload it
- Moreover, with in-memory replication we may have a strong enough guarantee for most purposes
  - Thus if durability costs enough to slow us down, we might opt for a weaker form of durability in order to gain better scalability and faster responses!

CS5412 Spring 2012 (Cloud Computing: Birman)

## This illustrates a challenge!

32

- Cloud systems just can't be approached in a one-size fits all manner
- For performance-intensive scalability scenarios we need to look closely at tradeoffs
  - Cost of stronger guarantee, versus
  - Cost of being faster but offering weaker guarantee
- If systems builders blindly opt for strong properties when not needed, we just incur other costs!
  - Amazon: Each 100ms delay reduces sales by 1%!

CS5412 Spring 2012 (Cloud Computing: Birman)

## Properties we might want

33

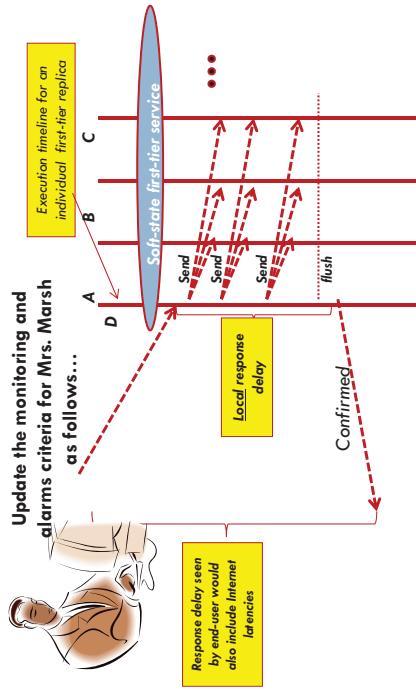
- Consistency: Updates in an agreed order
- Durability: Once accepted, won't be forgotten
- Real-time responsiveness: Replies with bounded delay
- Security: Only permit authorized actions by authenticated parties
- Privacy: Won't disclose personal data
- Resilience: Failures can't prevent the system from providing desired services
- Coordination: actions won't interfere with one another

CS5412 Spring 2012 (Cloud Computing: Birman)

# Fast response with consistency

## Does CAP apply deeper in the cloud?

34



This mixture of features gives us consistency, an in-memory replication guarantee ("amnesia freedom"), but not full durability

35

- The principle of wanting speed and scalability certainly is universal
- But many cloud services have strong consistency guarantees that we take for granted but depend on Marvin Theimer at Amazon explains:
  - Avoid costly guarantees that aren't even needed
  - But sometimes you just need to guarantee something
  - Then, be clever and engineer it to scale
  - And expect to revisit it each time you scale out 10x

CS5412 Spring 2012 (Cloud Computing: Birman)

## Cloud services and their properties

36

Service	Properties it guarantees
Memcached	No special guarantees
Google's GFS	File is current if locking is used
BigTable	Shared key-value store with many consistency properties
Dynamo	Amazon's shopping cart: eventual consistency
Databases	Snapshot isolation with log-based mirroring (a fancy form of the ACID guarantees)
MapReduce	Uses a "functional" computing model within which offers very strong guarantees
Zookeeper	Yahoo! file system with sophisticated properties
PNUTS	Yahoo! database system, sharded data, spectrum of consistency options
Chubby	Locking service... very strong guarantees

CS5412 Spring 2012 (Cloud Computing: Birman)

## Is there a conclusion to draw?

37

- One thing to notice about those services...
  - Most of them cost 10's or 100's of millions to create
  - Huge investment required to build strongly consistent and scalable and high performance solutions
  - Oracle's current parallel database: billions invested
  - CAP isn't about telling Oracle how to build a database product...
  - CAP is a warning to you that strong properties can easily lead to slow services
  - But thinking in terms of weak properties is often a successful strategy that yields a good solution and requires less effort

CS5412 Spring 2012 (Cloud Computing: Birman)

# Core problem?



38

- When can we safely sweep consistency under the rug?
- If we weaken a property in a safety critical context, something bad can happen!
- Amazon and eBay do well with weak guarantees because many applications just didn't need strong guarantees to start with!
- By embracing their weaker nature, we reduce synchronization and so get better response behavior
- But what happens when a wave of high assurance applications starts to transition to cloud-based models?

CS5412 Spring 2012 (Cloud Computing: Birman)

## Sommario

## Scopo del progetto

- ➊ Scopo del progetto
- ➋ Approccio architetture per la risoluzione delle problematiche affrontate
- ➌ Caso d'uso reale per verificare le soluzioni proposte
- ➍ Ambiente Cloud utilizzato per implementare l'architettura proposta
- ➎ Test effettuati
- ➏ Conclusioni
- ➐ Sviluppi futuri

## Scopo del progetto

*Le architetture delle piattaforme Cloud sono del tutto diverse da quelle convenzionali*

- Applicazioni ⇒ Fornite come servizi web
- Istanza ⇒ Entità operativa che fruisce i servizi
- Elasticità ⇒ Capacità di adattarsi (scaling) all'esigenza corrente
  - Otttenuta mediante la replicazione delle istanze

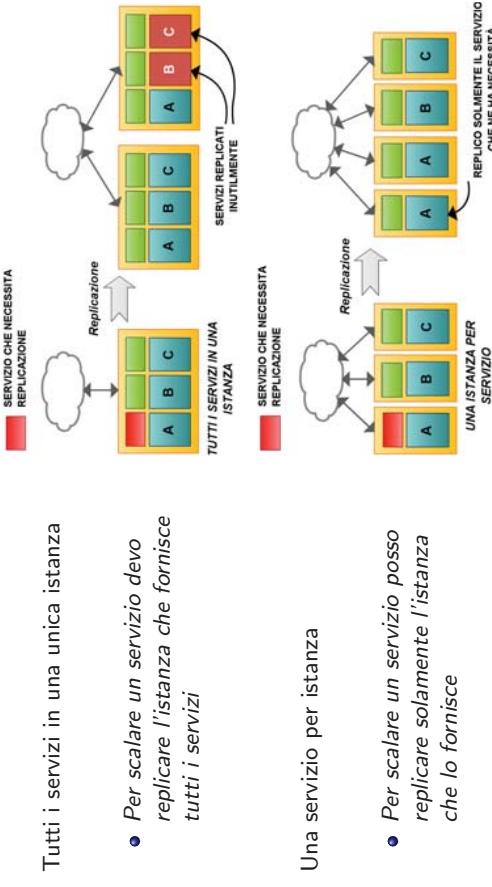


- Applicazioni progettate per piattaforme convenzionali sono inadeguate agli ambienti Cloud

*Quali considerazioni architettoniche sono necessarie?*

## Scomposizione dei servizi

### Necessario replicare solamente i componenti che ne hanno reale necessità

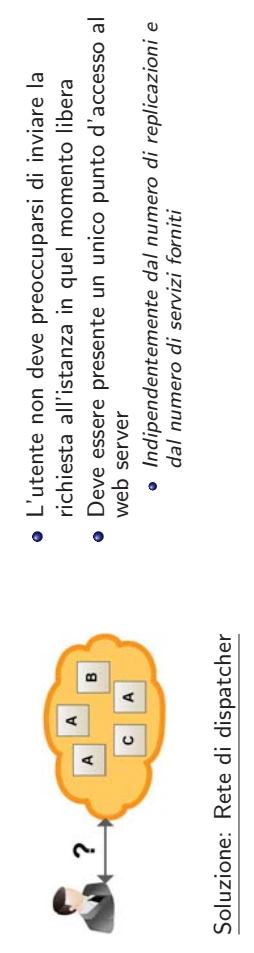


Federico Caccio – GeoServer nel Cloud

4/ 22

## Rete di Dispatcher

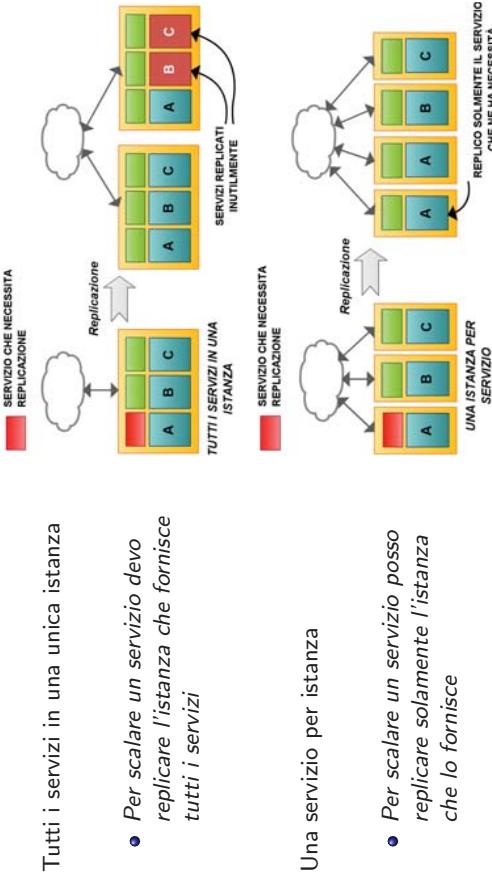
### La replicazione delle istanze che compongono il web server deve essere trasparente



5/ 22

## Configurazione dei servizi

### Necessario replicare solamente i componenti che ne hanno reale necessità



Federico Caccio – GeoServer nel Cloud

4/ 22

## Memorizzazione dei dati

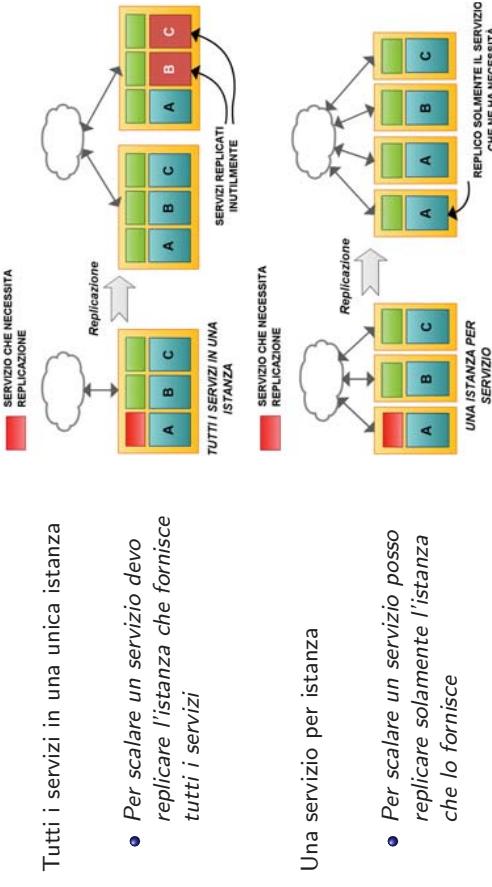


Federico Caccio – GeoServer nel Cloud

5/ 22

## Configurazione comune delle istanze

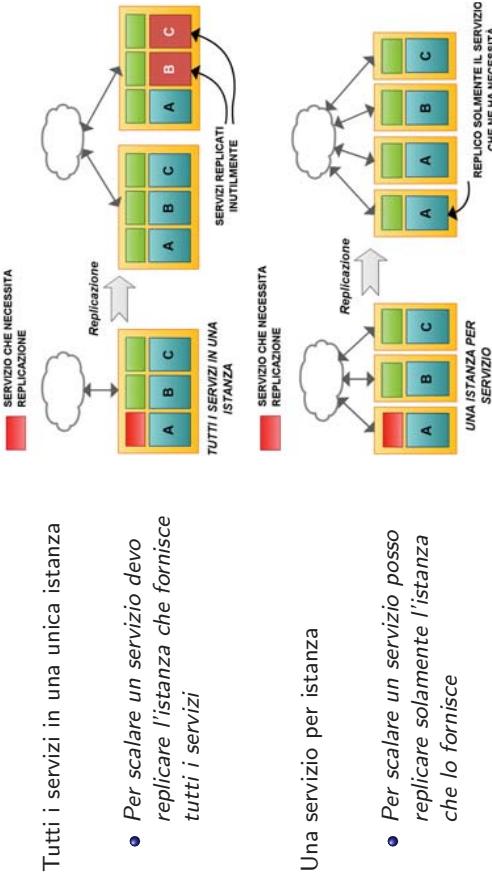
### Soluzione 1: Configurazione in uno spazio condiviso



Federico Caccio – GeoServer nel Cloud

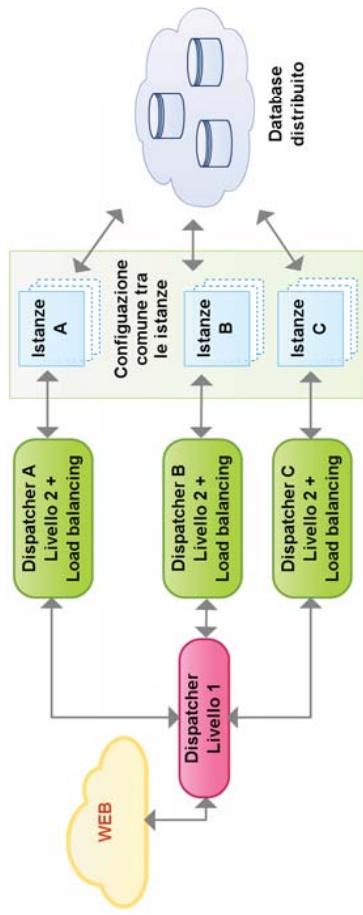
6/ 22

## La scelta dipende dall'architettura dei servizi e dalla piattaforma Cloud utilizzata



Federico Caccio – GeoServer nel Cloud

7/ 22



Scopo del progetto Approccio architettonico Caso d'uso Implementazione Test Conclusioni Sviluppi futuri  
Strumenti AWS utilizzati

Ambiente Cloud utilizzato: Amazon Web Services

- Rilevata nei Load Balancer di secondo livello

Scopo del progetto Approccio architettonico Caso d'uso Implementazione Test Conclusioni Sviluppi futuri  
Caso di studio concreto: GeoServer

Server web geospaziale che fornisce i servizi

- Web Coverage Service
- Web Feature Service
- Web Map Service

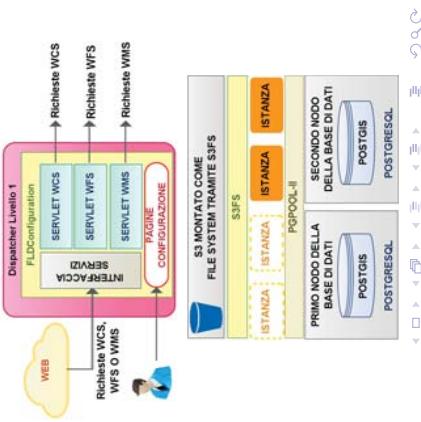
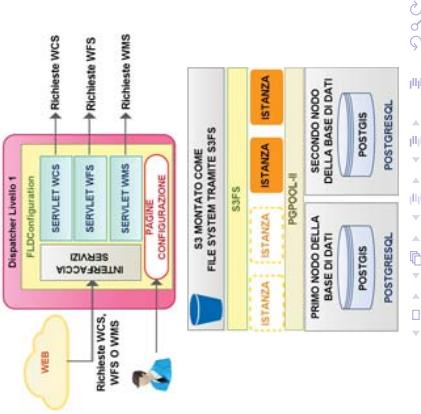
Servizi con differenti carichi computazionali

- Scomposizione per poter avviare istanze contenenti i singoli servizi

Scopo del progetto Approccio architettonico Caso d'uso Implementazione Test Conclusioni Sviluppi futuri  
Federico Caccio – GeoServer nel Cloud

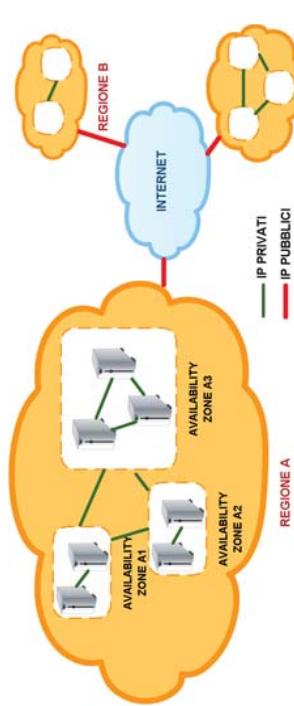
Strumenti esterni ad AWS utilizzati

- S3FS (Middleware supporto S3)
- Pgpool-II (Middleware di distribuzione)
- PostgreSQL (Data Base Management System)
- PostGIS (Estensione per dati geospatiali)



## Regioni e Availability Zone

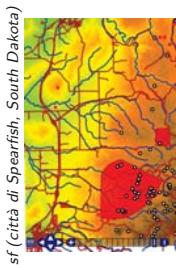
- Regione: Sedi su cui è collocato AWS nel mondo
- Availability Zone: Sotto suddivisione delle Regioni



Le Availability Zone consentono di implementare architetture affidabili  
impedendo la diffusione dei guasti

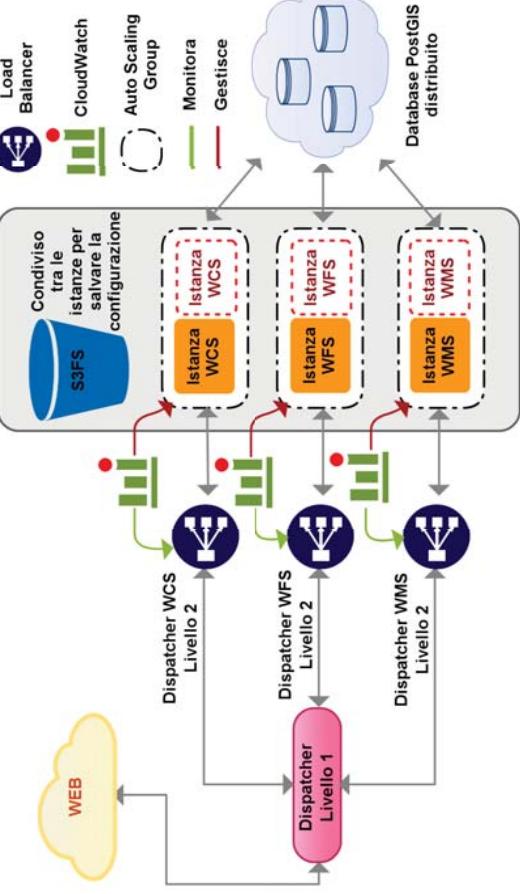
## Configurazione dei test

- Simulati utenti contemporanei che invocano i servizi
- 2 serie di dati geospatiali suddivisi in 2 workspace



- 2 configurazioni di distribuzione della base di dati
  - Partizionamento
  - Replicazione

## Visione d'insieme



## Availability Zone

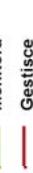
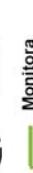
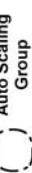
La distribuzione nelle Availability Zone per aumentare l'affidabilità introduce latenze?

- Simulati 5 utenti contemporanei

Richiesta	Diversa AZ	Lat. media (ms)	Throughput (rich/sec)	Lat. media (ms)	Stessa AZ	Throughput (rich/sec)
WMS.getMap (sf)	606	1.6180		656	1.6662	
WMS.getMap (tiger)	580	1.77667		556	1.6785	
WMS.getMap_multilayer_3.svg (sf)	943	1.7620		1001	1.6722	
WMS.getMap_multilayer_3.svg (tiger)	2412	1.7518		2610	1.6360	
Total	1134	6.9874		1205	6.5174	

Distribuire l'applicazione tra più Availability Zone non introduce latenze

## Scopo del progetto Approccio architettonico Caso d'uso Implementazione Test Conclusioni Sviluppi futuri



## Scopo del progetto Approccio architettonico Caso d'uso Implementazione Test Conclusioni Sviluppi futuri

## Scopo del progetto Approccio architettonico Caso d'uso Implementazione Test Conclusioni Sviluppi futuri

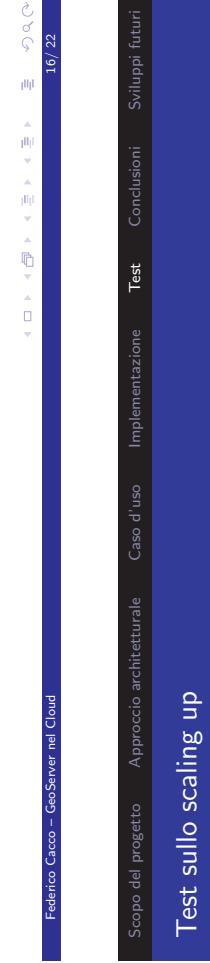
## Partizionato VS Replicato - 2 workspace

La distribuzione mediante partizionamento è più efficiente nel caso di richieste riferite a workspace diverse?

- Simulati 16 utenti contemporanei

Richiesta	Lat. media (ms)	DB partizionato Throughput (rich/sec)	Lat. media (ms)	DB replicato Throughput (rich/sec)
WMS-getMap (sf)	524	3.4943	818	2.7626
WMS-getMap (tiger)	577	3.4895	807	2.7604
WMS-getMap_multilayer_3.svg (sf)	975	3.4847	1322	2.7477
WMS-getMap_multilayer_3.svg (tiger)	2516	3.4379	2870	2.7571
Totali	1146	13.7263	1452	10.8429
		Totali	1248	12.5958
				1542
				10.1734

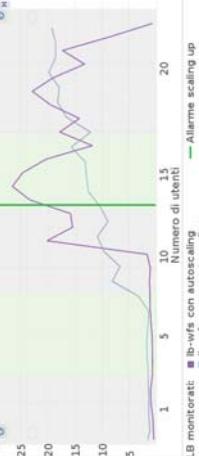
Con chiamate relative a dati in differenti workspace la distribuzione mediante partizionamento risulta più efficiente



Lo scaling up riduce la latenza quando essa diventa eccessiva?

Soglie scaling:

- Scaling up: lat. > 15 s (MAX)
- Scaling down: lat. < 1 s (MED)



Latenza ridotta dallo scaling up

## Partizionato VS Replicato - 1 workspace

La distribuzione mediante replicazione è più efficiente nel caso di richieste riferite sempre alla stessa workspace?

- Simulati 16 utenti contemporanei

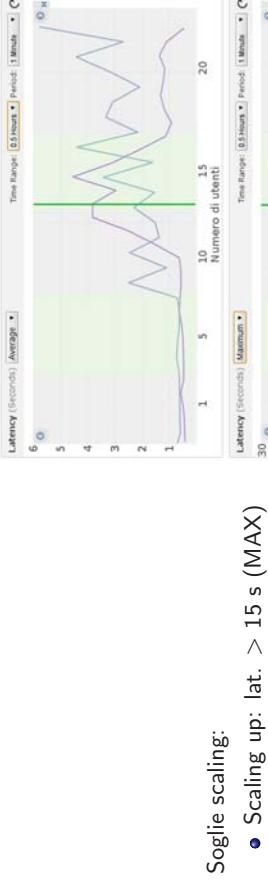
Richiesta	Lat. media (ms)	DB partizionato Throughput (rich/sec)	Lat. media (ms)	DB replicato Throughput (rich/sec)
WMS-getMap (tiger)	2.7626	3.2172	1021	2.6029
WMS-getMap_multilayer_3.svg (tiger)	3.1615	3.2277	2.5524	2.5952
WMS-getMap_multilayer_filter (tiger)	3.2058	925	2.5944	2.5944
Totali	3.1974	989	10.1734	10.1734

Distribuzione mediante partizionamento nuovamente più efficiente

Per piccole distribuzioni il middleware non introduce benefici



Vengono liberate le risorse quando non sono più necessarie?

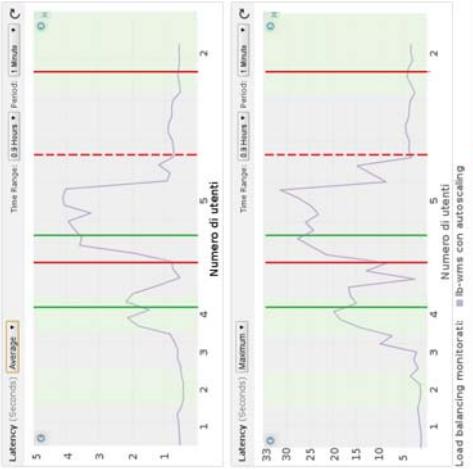


Soglie scaling:

- Scaling up: lat. > 15 s (MAX)
- Scaling down: lat. < 1 s (MED)

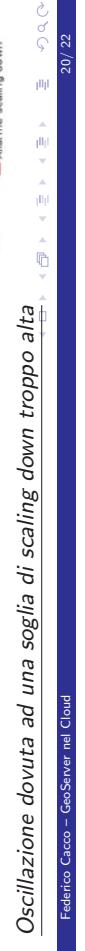
## Test sullo scaling down - oscillazione

Importante utilizzare soglie di scaling corrette!



Soglie scaling:

- Scaling up: lat. > 10 s (MAX)
- Scaling down 1: lat. < 1 s (MED)
- Scaling down 2: lat. < 0.6 s (MED)

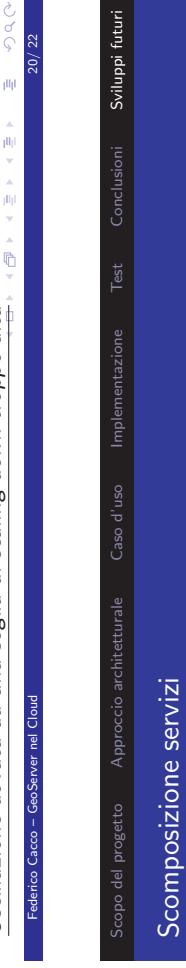


Federico Cucco – GeoServer nel Cloud

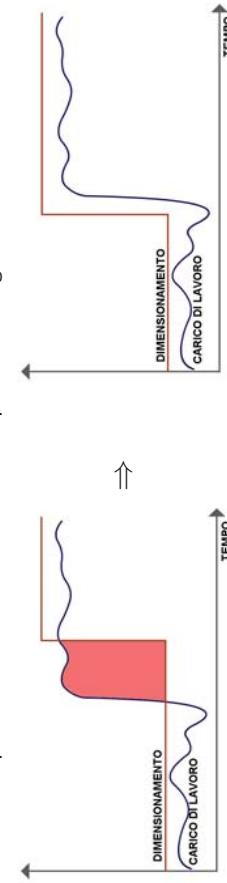
20/22

## Conclusioni

- La latenza si rivela una buona metrica per determinare se il dimensionamento è coerente con il carico di lavoro
- Rendere l'architettura più affidabile mediante una sua distribuzione su più Availability Zone non ne pregiudica le prestazioni
- Nel caso di richieste fatte a più workspace, la distribuzione mediante partizionamento si rileva più efficiente di quella realizzata mediante replicazione
- Lo scaling up è una operazione molto più onerosa, in termini di tempo e risorse, dello scaling down
- Difficile determinare soglie di scaling corrette



- Prendere i picchi di carico in modo da anticipare lo scaling dei servizi



- Testare la distribuzione mediante replicazione con scaling up di dimensioni maggiori