

# Sustainable Cloud Computing

<sup>1</sup>Amandeep Kaur, <sup>2</sup>Ranju Marwaha

<sup>1,2</sup>SSIET, Derabassi, Punjab, India

## Abstract

Cloud Computing is rising fast with its data centres growing at an unprecedented rate. However, this has come with concerns over privacy, efficiency at the expense of resilience, and environmental sustainability, because of the dependence on Cloud vendors such as Google, Amazon and Microsoft. Our response is an alternative model for the Cloud conceptualization, providing a paradigm for Clouds in the community, utilising networked personal computers for liberation from the centralised vendor model. Community Cloud Computing (C3) offers an alternative architecture, created by combining the Cloud with paradigms from Grid Computing, principles from Digital Ecosystems and sustainability from Green Computing, while remaining true to the original vision of the Internet. It is more technically challenging than Cloud Computing, having to deal with distributed computing issues, including heterogeneous nodes, varying quality of service, and additional security constraints. However, these are not insurmountable challenges, and with the need to retain control over our digital lives and the potential environmental consequences, it is a challenge we must pursue.

## Keywords

Cloud Computing, Community Cloud, Community Cloud Computing, Green Computing, Sustainability.

## I. Introduction

The recent development of Cloud Computing provides a compelling value proposition for organisations to out-source their Information and Communications Technology (ICT) infrastructure. However, there are growing concerns over the control ceded to large Cloud vendors, especially the lack of information privacy. Also, the data centres required for Cloud Computing are growing exponentially, creating an ever-increasing carbon footprint and therefore raising environmental concerns. The distributed resource provision from Grid Computing, distributed control from Digital Ecosystems, and sustainability from Green Computing, can remedy these concerns. So, Cloud Computing combined with these approaches would provide a compelling socio-technical conceptualisation for sustainable distributed computing, utilising the spare resources of networked personal computers collectively to provide the facilities of a virtual data centre and form a Community Cloud. Therefore, essentially reformulating the Internet to reflect its current uses and scale, while maintaining the original intentions for sustainability in the face of adversity. Including extra capabilities embedded into the infrastructure which would become as fundamental and invisible as moving packets is today.

## II. Cloud Computing

Cloud Computing is the use of Internet-based technologies for the provision of services, originating from the cloud as a metaphor for the Internet, based on depictions in computer network diagrams to abstract the complex infrastructure it conceals. It can also be seen as a commercial evolution of the academic-oriented Grid Computing, succeeding where Utility Computing struggled, while making greater use of the self-management advances of Autonomic Computing. It offers the illusion of infinite computing resources available on demand, with the elimination of upfront commitment from users, and payment for the use of computing

resources on a short-term basis as needed. Furthermore, it does not require the node providing a service to be present once its service is deployed. It is being promoted as the cutting-edge of scalable web application development, in which dynamically scalable and often virtualised resources are provided as a service over the Internet, with users having no knowledge of, expertise in, or control over the technology infrastructure of the Cloud supporting them. It currently has significant momentum in two extremes of the web development industry: the consumer web technology incumbents who have resource surpluses in their vast data centres, and various consumers and start-ups that do not have access to such computational resources. Cloud Computing conceptually incorporates Software-as-a-Service (SaaS) and other technologies with reliance on the Internet, providing common business applications online through web browsers to satisfy the computing needs of users, while the software and data are stored on the servers.

Fig. 1 shows the typical configuration of Cloud Computing at run-time when consumers visit an application served by the central Cloud, which is housed in one or more data centres. Green symbolises resource consumption, and yellow resource provision. The role of coordinator for resource provision is designated data centres. Green symbolises resource consumption and yellow




Fig. 1 : Cloud Computing

Typical configuration when consumers visit an application served by the central Cloud, which is housed in one or more resource provision. The role of coordinator for resource provision is designated by red and is centrally controlled by red, and is centrally controlled. Even if the central node is implemented as a distributed grid, which is the usual incarnation of a data centre, control is still centralised. Providers, who are the controllers, are usually companies with other web activities that require large computing

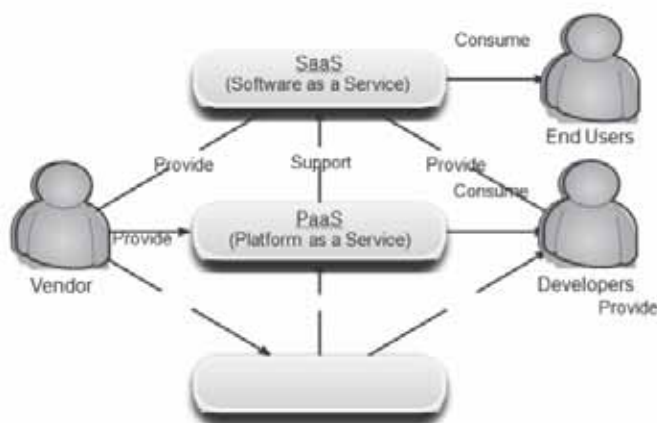


Fig. 2 : Abstractions of Cloud Computing

While there is a significant buzz around Cloud Computing, there is little clarity over which offerings qualify or their interrelation. The key to resolving this confusion is the realisation that the various offerings fall into different levels of abstraction, aimed at different market segments. Resources, and in their efforts to scale their primary businesses have gained considerable expertise and hardware. For them, Cloud Computing is a way to resell these as a new product while expanding into a new market. Consumers include everyday users, Small and Medium sized Enterprises (SMEs), and ambitious start-ups whose innovation potentially threatens the incumbent providers.

### A. Layers of Abstraction

While there is a significant buzz around Cloud Computing, there is little clarity over which offerings qualify or their interrelation. The key to resolving this confusion is the realisation that the various offerings fall into different levels of abstraction, as shown in Fig. 2, aimed at different market segments.

#### 1. Infrastructure-as-a-Service (IaaS)

At the most basic level of Cloud Computing offerings, there are providers such as Amazon and Mosso, who provide machine instances to developers. These instances essentially behave like dedicated servers that are controlled by the developers, who therefore have full responsibility for their operation. So, once a machine reaches its performance limits, the developers have to manually instantiate another machine and scale their application out to it. This service is intended for developers who can write arbitrary software on top of the infrastructure with only small compromises in their development methodology.

#### 2. Platform-as-a-Service (PaaS)

One level of abstraction above, services like Google App Engine provide a programming environment that abstracts machine instances and other technical details from developers. The programs are executed over data centres, not concerning the developers with matters of allocation. In exchange for this, the developers have to handle some constraints that the environment imposes on their application design, for example the use of key-value stores instead of relational databases.

#### 3. Software-as-a-Service (SaaS)

At the consumer-facing level are the most popular examples of Cloud Computing, with well-defined applications offering users online resources and storage. This differentiates SaaS from traditional websites or web applications which do not interface

with user information (e.g. documents) or do so in a limited manner. Popular examples include Microsoft's (Windows Live) Hotmail, office suites such as Google Docs and Zoho, and online business software such as Salesforce.com. To better understand Cloud Computing we can categorise the roles of the various actors. The vendor as resource provider has already been discussed. The application developers utilise the resources provided, building services for the end users. This separation of roles helps define the stakeholders and their differing interests. However, actors can take on multiple roles, with vendors also developing services for the end users, or developers utilising the services of others to build their own services. Yet, within each Cloud the role of provider, and therefore controller, can only be occupied by the vendor providing the Cloud.

### B. Concerns

The Cloud Computing model is not without concerns, as others have noted and we consider the following as primary:

#### 1. Failure of Monocultures

The uptime of Cloud Computing based solutions is an advantage, when compared to businesses running their own infrastructure, but often overlooked is the co-occurrence of downtime in vendor-driven monocultures. The use of globally decentralised data centres for vendor Clouds minimises failure, aiding its adoption. However, when a cloud fails, there is a cascade effect crippling all organisations dependent on that Cloud, and all those dependent upon them. This was illustrated by the Amazon (S3) Cloud outage, which disabled several other dependent businesses. So, failures are now system-wide, instead of being partial or localised. Therefore, the efficiencies gained from centralising infrastructure for Cloud Computing are increasingly at the expense of the Internet's resilience.

#### 2. Convenience vs Control

The growing popularity of Cloud Computing comes from its convenience, but also brings vendor control, an issue of ever-increasing concern. For example, Google Apps for in-house e-mail typically provides higher uptime, but its failure highlights the issue of lock-in that comes from depending on vendor Clouds. The even greater concern is the loss of information privacy, with vendors having full access to the resources stored on their Clouds. So much so the British government is considering a 'G Cloud' for government business applications. In particularly sensitive cases of SMEs and start-ups, the provider-consumer relationship that Cloud Computing fosters between the owners of resources and their users could potentially be detrimental, as there is a potential conflict of interest for the providers. They profit by providing resources to up-and-coming players, but also wish to maintain dominant positions in their consumer-facing industries.

#### 3. Environmental Impact

The other major concern is the ever-increasing carbon footprint from the exponential growth of the data centres required for Cloud Computing. With the industry expected to exceed the airline industry by 2020 raising sustainability concerns. The industry is being motivated to address the problem by legislation, the operational limit of power grids (being unable to power anymore servers in their data centres) and the potential financial benefits of increased efficiency. Their primary solution is the use of virtualisation to maximise resource utilisation, but the problem remains. While these issues are endemic to Cloud Computing, they are not flaws in the Cloud conceptualisation, but the vendor

provision and implementation of Clouds. There are attempts to address some of these concerns, such as a portability layer between vendor Clouds to avoid lock-in. However, this will not alleviate issues such as inter-Cloud latency. An open source implementation of the Amazon (EC2) Cloud, called Eucalyptus, allows a data centre to execute code compatible with Amazon's Cloud. Allowing for the creation of private internal Clouds, avoiding vendor lock-in and providing information privacy, but only for those with their own data centre and so is not really Cloud Computing (which by definition is to avoid owning data centres). Therefore, vendor Clouds remain synonymous with Cloud Computing. Our response is an alternative model for the Cloud conceptualisation, created by combining the Cloud with paradigms from Grid Computing, principles from Digital Ecosystems, and sustainability from Green Computing, while remaining true to the original vision of the Internet.

### III. Grid Computing : Distributing Provision

Grid Computing is a form of distributed computing in which a virtual super computer is composed from a cluster of networked, loosely coupled computers, acting in concert to perform very large tasks. It has been applied to computationally intensive scientific, mathematical, and academic problems through volunteer computing, and used in commercial enterprise for such diverse applications as drug discovery, economic forecasting, seismic analysis, and back-office processing to support e-commerce and web services.




Fig. 3 : Grid Computing

Typical configuration in which resource provision is managed by a group of distributed nodes. Green symbolises resource consumption, and yellow resource provision. The role of coordinator for resource provision is designated by red, and is centrally controlled. What distinguishes Grid Computing from cluster computing is being more loosely coupled, heterogeneous, and geographically dispersed. Also, grids are often constructed with general-purpose grid software libraries and middleware, dividing and apportioning pieces of a program to potentially thousands of computers. However, what distinguishes Cloud Computing from Grid Computing is being web-centric, despite some of its definitions being conceptually similar (such as computing resources being consumed as electricity is from power grids).

### IV. Digital Ecosystems: Distributing Control

Digital Ecosystems are distributed adaptive open socio-technical systems, with properties of self-organisation, scalability and sustainability, inspired by natural ecosystems. Emerging as a novel approach to the catalysis of sustainable regional development driven by SMEs. Aiming to help local economic actors become active players in globalisation, valorising their local culture and vocations, and enabling them to interact and create value networks at the global level. Increasingly this approach, dubbed glocalisation, is being considered a successful strategy of globalisation that preserves regional growth and identity and has been embraced by the mayors and decision-makers of thousands of municipalities [56]. The community focused on the deployment of Digital Ecosystems, Regions for Digital Ecosystems Network (REDEN), is supported by projects such as the Digital Ecosystems Network of regions for Dissemination and Knowledge Deployment (DEN4DEK). This thematic network that aims to share experiences and disseminate knowledge to let regions effectively deploy of Digital Ecosystems at all levels (economic, social, technical and political) to produce real impacts in the economic activities of European regions through the improvement of SME business environments.

In a traditional market-based economy, made up of sellers and buyers, the parties exchange property, while in a new network-based economy, made up of servers and clients, the parties share access to services and experiences. Digital Ecosystems aim to support network-based economies reliant on next-generation ICT that will extend the Service-Oriented Architecture (SOA) concept [59] with the automatic combining of available and applicable services in a scalable architecture, to meet business user requests for applications that facilitate business processes. Digital Ecosystems research is yet to consider scalable resource provision, and therefore risks being subsumed into vendor Clouds at the infrastructure level, while striving for decentralisation at the service level. So, the realisation of their vision requires a form of Cloud Computing, but with their principle of community-based infrastructure where individual users share ownership.

### V. Green Computing: Growing Sustainably

Green Computing is the efficient use of computing resources, with the primary objective being to account for the triple bottom line, an expanded spectrum of values and criteria for measuring organizational (and societal) success. Given computing systems existed before concern over their environmental impact, it has generally been implemented retroactively, but is now being considered at the development phase. It is systemic in nature, because ever-increasingly sophisticated modern computer systems rely upon people, networks and hardware. So, the elements of a green solution may comprise items such as end user satisfaction, management restructuring, regulatory compliance, disposal of electronic waste, telecommuting, virtualisation of server resources, energy use, thin client solutions and return on investment. One of the greatest environmental concerns of the industry is their data centres, which have increased in number over time as business demands have increased, with facilities housing a rising amount of evermore powerful equipment. As data centres run into limits related to power, cooling and space, their ever-increasing operation has created a noticeable impact on power grids. To the extent that data centre efficiency has become an important global issue, leading to the creation of the Green Grid, an international non-profit organisation mandating an increase in the energy efficiency of data centres. Their approach, virtualisation, has improved efficiency, but is optimising a flawed model that does not consider

the whole system, where resource provision is disconnected from resource consumption. For example, competing vendors must host significant redundancy in their data centres to manage usage spikes and maintain the illusion of infinite resources. So, we would argue that an alternative more systemic approach is required, where resource consumption and provision are connected, to minimise the environmental impact and allow sustainable growth.

## VI. Community Cloud

C3 arises from concerns over Cloud Computing, specifically control by vendors and lack of environmental sustainability. The Community Cloud aspires to combine distributed resource provision from Grid Computing, distributed control from Digital Ecosystems and sustainability from Green Computing, with the use cases of Cloud Computing, while making greater use of self-management advances from Autonomic Computing. Replacing vendor Clouds by shaping the underutilised resources of user machines to form a Community Cloud, with nodes potentially fulfilling all roles, consumer, producer, and most importantly coordinator.

### A. Conceptualisation

The conceptualisation of the Community Cloud draws upon Cloud Computing, Grid Computing, Digital Ecosystems, Green Computing and Autonomic Computing. A paradigm for Cloud Computing in the community, without dependence on Cloud vendors, such as Google, Amazon, or Microsoft.

#### 1. Openness

Removing dependence on vendors makes the Community Cloud the open equivalent to vendor Clouds, and therefore identifies a new dimension in the open versus proprietary struggle that has emerged in code, standards and data, but has yet to be expressed in the realm of hosted services.

#### 2. Community

The Community Cloud is as much a social structure as a technology paradigm, because of the community ownership of the infrastructure. Carrying with it a degree of economic scalability, without which there would be diminished competition and potential stifling of innovation as risked in vendor Clouds.

#### 3. Individual Autonomy

In the Community Cloud, nodes have their own utility functions in contrast with data centres, in which dedicated machines execute software as instructed. So, with nodes expected to act in their own self-interest, centralised control would be impractical, as with consumer electronics like game consoles. Attempts to control user machines counter to their self-interest results in cracked systems, from black market hardware modifications and arms races over hacking and securing the software (routinely lost by the vendors) [66]. In the Community Cloud, where no concrete vendors exist, it is even more important to avoid antagonising the users, instead embracing their self interest and harnessing it for the benefit of the community with measures such as a community currency.

#### 4. Identity

In the Community Cloud each user would inherently possess a unique identity, which combined with the structure of the Community Cloud should lead to an inversion of the currently predominant membership model. So, instead of users registering for each website (or service) anew, they could simply add the website to their identity and grant access. Allowing users to have

multiple services connected to their identity, instead of creating new identities for each service. This relationship is reminiscent of recent application platforms, such as Facebook's f8 and Apple's App Store, but decentralised in nature and so free from vendor control. Also, allowing for the reuse of the connections between users, akin to Google's Friend Connect, instead of reestablishing them for each new application.

### 5. Graceful Failures

The Community Cloud is not owned or controlled by any one organisation, and therefore not dependent on the lifespan or failure of any one organisation. It therefore ought to be robust and resilient to failure, and immune to the system-wide cascade failures of vendor Clouds, because of the diversity of its supporting nodes. When occasionally failing doing so gracefully, non-destructively, and with minimal downtime, as the unaffected nodes mobilise to compensate for the failure.

### 6. Convenience and Control

The Community Cloud, unlike vendor Clouds, has no inherent conflict between convenience and control, resulting from its community ownership providing distributed control, which would be more democratic. However, whether the Community Cloud can provide technically quality equivalent or superior to its centralised counterparts is an issue that will require further research.

### 7. Community Currency

The Community Cloud would require its own currency to support the sharing of resources, a community currency, which in economics is a medium (currency), not backed by a central authority (e.g. national government), for exchanging goods and services within a community. It does not need to be restricted geographically, despite sometimes being called a local currency. An example is the Fureai kippu system in Japan, which issues credits in exchange for assistance to senior citizens. Family members living far from their parents can earn credits by offering assistance to the elderly in their local community, which can then be transferred to their parents and redeemed by them for local assistance.

### 8. Quality of Service

Ensuring acceptable quality of service (QoS) in a heterogeneous system will be a challenge. Not least because achieving and maintaining the different aspects of QoS will require reaching critical mass in the participating nodes and available services. Thankfully, the community currency could support long-term promises by resource providers and allow the higher quality providers, through market forces, to command a higher price for their service provision. Interestingly, the Community Cloud could provide a better QoS than vendor Clouds, utilising time-based and geographical variations advantageously in the dynamic scaling of resource provision.

### 9. Environmental Sustainability

We expect the Community Cloud to have a smaller carbon footprint than vendor Clouds, on the assumption that making use of underutilised user machines requires less energy than the dedicated data centres required for vendor Clouds. The server farms within data centres are an intensive form of computing resource provision, while the Community Cloud is more organic, growing and shrinking in a symbiotic relationship to support the demands of the community, which in turn supports it.



## 10. Service Composition

The great promise of service-oriented computing is that the marginal cost of creating the  $n$ -th application will be virtually zero, as all the software required already exists to satisfy the requirements of other applications. Only their composition and orchestration are required to produce a new application. Within vendor Clouds it is possible to make services that expose themselves for composition and compose these services, allowing the hosting of a complete service-oriented architecture. However, current service composition technologies have not gained widespread adoption. Digital Ecosystems advocate service composability to avoid centralised control by large service providers, because easy service composition allows coalitions of SMEs to compete simply by composing simpler services into more complex services that only large enterprises would otherwise be able to deliver. So, we should extend decentralisation beyond resource provision and up to the service layer, to enable service composition within the Community Cloud.

## B. Architecture

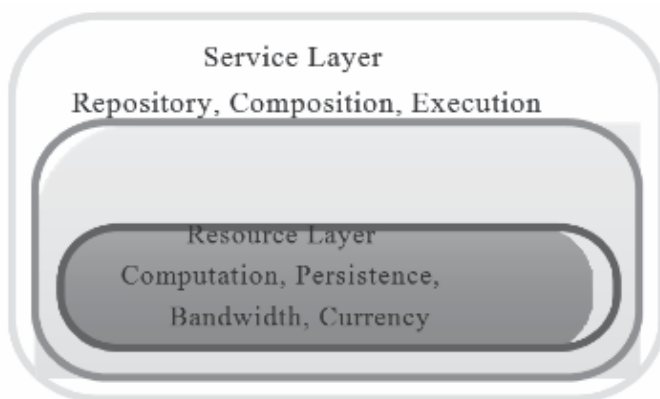


Fig. 5 : Community Cloud Computing

An architecture in which the most fundamental layer deals with distributing coordination. One layer above, resource provision and consumption are arranged on top of the coordination framework. Finally, the service layer is where resources are combined into end-user accessible services, to then themselves be composed into higher-level services. The method of materializing the Community Cloud is the distribution of its server functionality amongst a population of nodes provided by user machines, shaping their underutilised resources into a virtual data centre. While straightforward in principle, it poses challenges on many different levels. So, an architecture for C3 can be divided into three layers, dealing with these challenges iteratively. The most fundamental layer deals with distributing coordination, which is taken for granted in homogeneous data centres where good connectivity, constant presence and centralised infrastructure can be assumed. One layer above, resource provision and consumption are arranged on top of the coordination framework. Easy in the homogeneous grid of a data centre where all nodes have the same interests, but more challenging in a distributed heterogeneous environment. Finally, the service layer is where resources are combined into end-user accessible services, to then themselves be composed into higher-level services.

## 1. Coordination Layer

To achieve coordination, the nodes need to be deployed as isolated virtual machines, forming a fully distributed P2P6 network that can provide support for distributed identity, trust, and transactions.

## A. Virtual Machines (VMs)

Executing arbitrary code in the machine of a resource-providing user would require a sandbox<sup>7</sup> for the guest code, a VM<sup>8</sup> to protect the host. The role of the VM is to make system resources safely available to the Community Cloud, upon which Cloud processes could be run safely (without danger to the host machine). Fortunately, feasibility has been proven with heavyweight VMs such as the Java Virtual Machine, lightweight JavaScript VMs present in most modern web browsers, and new approaches such as Google's Native Client. Furthermore, the age of multi-core processors has resulted in unused and underutilised cores being commonplace in modern personal computers, which lend themselves well to the deployment and background execution of Community Cloud facing VMs. Regarding deployment, users would be required to maintain an active browser window or tab, or install a dedicated application. While the first would not require installation privileges, the later would with the benefit of greater functionality. However, more likely a hybrid of both would occur, facilitating the availability and advantages of each in different scenarios.

## B. Distributed Identity

In distributed systems with variable node reliability, historical context is logically required to have certainty of node interactions. Fundamental to this context is the ability to identify nodes and therefore reference previous interactions. However, current identification schemes have identity providers controlling provision. Such as in the DNS, which while nominally distributed, remains under centralised control both technologically and organisationally, permitting numerous distortions in the network. Including domain squatting<sup>11</sup>, abuses by domain registrars, subjection to political control and risks to the infrastructure being compromised. Identity in the Community Cloud has to arise naturally from the structure of the network, based on the relation of nodes to each other, so that it can scale and expand without centralised control. We can utilise the property that a large enough identifier-space is unlikely to suffer collisions. For example, the Git distributed version control system assigns a universal identifier to each new commission, without coordination with other repositories. Analogously, assuming each node independently produces a private-public key pair, the probability of public key collision is negligible. Also, from the human identification of nodes we can utilise the property that each node, despite formal identity, possesses a unique position in the network, i.e. set of connections to other nodes. So, combining these two properties provides reasonable certainty for a distributed identity model where universal identification can be accomplished without centralised mediation, but this is still an active area of research.

## C. Networking

At this level, nodes should be inter-connected to form a P2P network. Engineered to provide high resilience while avoiding single points of control and failure, which would make decentralised super-peer based control mechanisms insufficient. Newer P2P designs offer sufficient guarantees of distribution, immunity to super-peer failure, and resistance to enforced control. For example, in the Distributed Virtual Super-Peer (DVSP) model a collection of peers logically combine to form a virtual super-peer, which dynamically changes over time to facilitate fluctuating demands.

## D. Distributed Transactions

A key element of distributed coordination is the ability of nodes to jointly participate in transactions that influence their individual

state. Appropriately annotated business processes can be executed over a distributed network with a transactional model maintaining the ACID properties on behalf of the initiator. Newer transaction models maintain these properties while increasing efficiency and concurrency. Other directions of research include relaxing these properties to maximise concurrency. Others still, focus on distributing the coordination of transactions. A feature vital for C3, as distributed transaction capabilities are fundamental to permitting multi-party service composition without centralised mediation.

## 2. Resource layer

With the networking infrastructure now in place, we can consider the first consumer-facing uses for the virtual data centre of the Community Cloud. Offering the usage experience of Cloud Computing on the PaaS layer and above, because Cloud Computing is about using resources from the Cloud. So, Utility Computing scenarios, such as access to raw storage and computation, should be available at the PaaS layer. Access to these abstract resources for service deployment would then provide the SaaS layer.

### A. Distributed Computation

The field has a successful history of centrally controlled incarnations. However, C3 should also take inspiration from Grid Computing and Digital Ecosystems to provide distributed coordination of the computational capabilities that nodes offer to the Community Cloud.

### B. Distributed Persistence

The Community Cloud would naturally require storage on its participating nodes, taking advantage of the ever-increasing surplus on most personal computers. However, the method of information storage in the Community Cloud is an issue with multiple aspects. First, information can be file-based or structured. Second, while constant and instant availability can be crucial, there are scenarios in which recall times can be relaxed. Such varying requirements call for a combination of approaches, including distributed storage, distributed databases and key-value stores. Information privacy in the Community Cloud should be provided by the encryption of user information when on remote nodes, only being unencrypted when accessed by the user, allowing for the secure and distributed storage of information.

### C. Bandwidth Management

The Community Cloud would probably require more bandwidth at the user nodes than vendor Clouds, but can take advantage of the ever-increasing bandwidth and deployment of broadband. Also, P2P protocols such as BitTorrent make the distribution of information over networks much less bandwidth-intensive for content providers, accomplished by using the downloading peers as repeaters of the information they receive. C3 should adopt such approaches to ensure the efficient use of available network bandwidth, avoiding fluctuations and sudden rises in demand (e.g. the Slashdot effect) burdening parts of the network.

### D. Community Currency

An important theme in the Community Cloud is that of nodes being contributors as well as consumers, which would require a community currency (redeemable against resources in the community) to reward users for offering resources. This would also allow traditional Cloud vendors to participate by offering their resources to the Community Cloud to gather considerable community currency, which they can then monetise against

participants running a community currency deficit (i.e. contributing less than they consume). The relative cost of resources (storage, computation, bandwidth) should fluctuate based on market demand, not least because of the impracticality of predicting or hard-coding such ratios. So, a node of the network would gather community currency by performing tasks for the community, which its user could then use to access resources of the Community Cloud.

## E. Resource Repository

Given that each node providing resources has a different location in the network and quality characteristics, a distributed resource repository would be required that could respond to queries for resources according to desired performance profiles. Such a query would have to consider historical performance, current availability, projected cost and geographical distribution of the nodes to be returned. A constraint optimisation problem, the results returned would be a set of nodes that fit the required profile, proportionally to the availability of suitable nodes.

## 3. Service Layer

Cloud Computing represents a new era for service-oriented architectures, making services explicitly dependent on other resource providers instead of building on self-sufficient resource locations. C3 makes this more explicit, breaking down the stand-alone service paradigm, with any service by default being composed of resources contributed by multiple participants. So, the following sections define the core infrastructural services that the Community Cloud would need to provide.

### A. Distributed Service Repository (DSR)

The service repository of the Community Cloud must provide persistence, as with traditional service repositories, for the pointers to services and their semantic descriptions. To support the absence of service-producing nodes during service execution, there must also be persistence of the executable code of services. Naturally, the implementation of a distributed service repository is made easier by the availability of the distributed storage infrastructure of the Community Cloud.

### B. Service Deployment and Execution

When a service is required, but is not currently instantiated on a suitable node, a copy should be retrieved from the DSR and instantiated as necessary, allowing for flexible responsiveness and resilience to unpredictable traffic spikes. As nodes are opportunistically interested in executing services to gather community currency for their users, so developers should note the resource cost of their services in their descriptions, allowing for pre-execution resource budgeting, and post-execution community currency payments. Being in a developer's own interest to mark resource costs correctly, because over-budgeting would burden their users and under-budgeting would cause premature service termination. Additionally, developers could add a subsidy to promote their services. Remote service execution would need to be secured against potentially compromised nodes, perhaps through encrypted processing schemes. Otherwise, such nodes while unable to access a complete traffic log of the services they execute, could potentially access the business logic; and we would be replacing the vendor introspection problem, with an anyone introspection problem. Since delivering a service over large distances in the network comes at a potentially high cost, the lack of a central well-connected server calls for a fundamental paradigm shift, from pull-oriented approaches to hybrid push/pull-

oriented approaches. So, instead of the pull-oriented approach of supplying services only upon request, service provision should also follow a push-oriented approach of preemptive deployment to strategically suitable nodes, including modifying their deployment profile based on the traffic patterns they face at run time.

### C. Programming Paradigm

A key innovation of Cloud Computing in its PaaS incarnation, is the offering of a well-specified context (programming paradigm) within which the services should be executed. The programming paradigm that produces these services is also important to C3, because it forms a contract between the service developers and resource providers. The current state-of-the-art requires manipulation of source code in which each line is context dependent, and so a single intended change may necessitate significant alterations at different locations in the codebase. A paradigm shift to declarative generative programming would be greatly beneficial, avoiding the need to manually manage cascading changes to the codebase. As the requirements behind a service would be made explicit and executable, and being human readable could therefore be manipulated directly as stand-alone artifacts. Additionally, barriers to service composition would be significantly decreased, beneficial to C3 and beyond.

### D. Distributed Innovation

When considering the Community Cloud over time, current software distribution models would cause problems. Should the infrastructure be dependent on a single provider for updates, they would become a single point of control, and possibly failure. Entrusting a single provider with the power to control the evolution of the architecture, even if they are considered benevolent, risks the development goals becoming misaligned with the community. Therefore, the Community Cloud should follow an evolutionary software distribution model. Extending an already-growing trend of using distributed code repositories such as Git and Mercurial over centralised code repositories such as Subversion and CVS. So, modifications to services, including infrastructural ones, should be distributed locally to migrate over the Community Cloud from where they are deployed, making use of the existing relationships between users. Users or their nodes (by default) could even choose to follow the updates that other trusted peers adopt. Therefore, new versions of a service would compete with older versions, and where superior (fitter) would distribute more widely, spreading further across the Community Cloud. So, updates to services would permeate through the network, in a distributed but regulated manner. We could even consider the updates to services, as the release of patches (modifications), allowing for frequent, smaller and iterative releases more akin to an evolutionary software distribution model. Potential speciation (branching) would encourage developers to co-ordinate their releases and ensure their patches are viable across different branches. Obviously, the ability to undo patches and step back through versions of infrastructural services would be necessary to maintain the Community Cloud. Still, without a more granular approach to conflict resolution from different patching sources, poor developer relations could risk fragmentation of the codebase and network. So, an alternative non-centralised software innovation model would be required, such as the declarative generative programming paradigm mentioned.

### VII. Conclusions

We have presented the Community Cloud as an alternative to Cloud Computing, created from blending its usage scenarios with paradigms from Grid Computing, principles from Digital

Ecosystems, self-management from Autonomic Computing, and sustainability from Green Computing. So, C3 utilises the spare resources of networked personal computers to provide the facilities of data centres, such that the community provides the computing power for the Cloud they wish to use. A socio-technical conceptualisation for sustainable distributed computing. While the Open Cloud Manifesto is well intentioned, its promotion of open standards for vendor Cloud interoperability has proved difficult. We believe it will continue to prove difficult until a viable alternative, such as C3, is developed. Furthermore, we hope that the Community Cloud will encourage innovation in vendor Clouds, forming a relationship analogous to the creative tension between open source and proprietary software. In the future we will continue to refine the various elements of C3, such as suitable mechanisms for a community currency, distributed alternatives to DNS, DVSPs, Restful Clouds, declarative generative programming paradigms, distributed innovation, and the environmental impact of the Community Cloud relative to vendor Clouds.

### References

- [1] M. Haynie, "Enterprise cloud services: Deriving business value from Cloud Computing," Micro Focus, Tech. Rep., 2009.
- [2] R. L. Peter Lucas, Joseph Ballay, "The wrong cloud," MAYA Design, Inc, Tech. Rep., 2009. [Online] Available: [http://www.maya.com/file\\_download/126/The%20Wrong%20Cloud.pdf](http://www.maya.com/file_download/126/The%20Wrong%20Cloud.pdf)
- [3] M. Armbrust, A. Fox, R. Griffith, A. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, M. Zaharia, "Above the Clouds: A Berkeley view of Cloud Computing," University of California, Berkeley, 2009. [Online] Available: <http://d1smfj0g31qzek.cloudfront.net/abovetheclouds.pdf>
- [4] J. Hayes, "Cred or croak?" IET Knowledge Network, Tech. Rep., 2008. [Online] Available: <http://kn.theiet.org/magazine/issues/0820/cred-croak-0820.cfm?SaveToPDF>
- [5] P. McKenna, "Can we stop the internet destroying our planet?" New Scientist, vol. 197, no. 2637, pp. 20–21, 2008.
- [6] J. Kaplan, W. Forrest, N. Kindler, "Revolutionizing data center energy efficiency," McKinsey & Company, Tech. Rep., 2008. [Online] Available: [http://www.mckinsey.com/client-service/bto/poointofview/pdf/Revolutionizing Data Center Efficiency.pdf](http://www.mckinsey.com/client-service/bto/poointofview/pdf/Revolutionizing%20Data%20Center%20Efficiency.pdf)
- [7] B. Leiner, V. Cerf, D. Clark, R. Kahn, L. Kleinrock, D. Lynch, J. Postel, L. Roberts, S. Wolff, "A brief history of the internet," Institute for Information Systems and Computer Media, Tech. Rep., 2001. [Online] Available: [http://www.iicm.tugraz.at/thesis/cguetl/diss/literatur/Kapitel02/References/Leiner et al. 2000/brief.html?timestamp=1197467969844](http://www.iicm.tugraz.at/thesis/cguetl/diss/literatur/Kapitel02/References/Leiner%20et%20al.%202000/brief.html?timestamp=1197467969844)
- [8] J. Scanlon, B. Wieners, "The internet cloud," The Industry Standard, Tech. Rep., 1999. [Online] Available: <http://www.thestandard.com/article/0,1902,5466,00.html>
- [9] I. Foster, Y. Zhao, I. Raicu, S. Lu, "Cloud Computing and Grid Computing 360-degree compared," in Grid Computing Environments Workshop, 2008, pp. 1–10.
- [10] T. Foremski, "Sun services CTO says utility computing acceptance is slow going," ZDNet, CBS Interactive, Tech. Rep., 2006. [Online] Available: <http://blogs.zdnet.com/Foremski/?p=33>
- [11] A. Orlowski, "The Cell chip - how will MS and Intel face the music?" The Register, Tech. Rep., 2005. [Online] Available: [http://www.theregister.co.uk/2005/02/03/cell\\_analysis\\_part two/](http://www.theregister.co.uk/2005/02/03/cell_analysis_part_two/)

- [12] J. Kephart, D. Chess, I. Center, N. Hawthorne, "The vision of autonomic computing," *Computer*, vol. 36, no. 1, pp. 41–50, 2003.
- [13] G. Gruman, E. Knorr, "What Cloud Computing really means," InfoWorld Inc., Tech. Rep., 2008. [Online] Available: <http://www.infoworld.com/article/08/04/07/15FE-cloud-computing-reality-1.html>
- [14] Gartner, "Cloud Computing will be as influential as e-business," Gartner, Tech. Rep., 2008. [Online] Available: <http://www.gartner.com/it/page.jsp?id=707508>