



**INTERNATIONAL JOURNAL
OF TECHNOVATION AND
BUSINESS INSIGHTS**

**Volume-1, Issue-1, July 2025
Inaugural Edition**

PUBLISHED BY: D3 PUBLISHERS



Cloud Computing: A State-of-the-Art Review and Emerging Research Challenges

Dr.SK YAKOOB, Associate Professor & HOD,

Department of CSE,

SAI SPURTHY INSTITUTE OF TECHNOLOGY(AUTONOMOUS), KHAMMAM.

ABSTRACT: Cloud computing is transforming online data storage and interchange for enterprises of all sizes with its adaptable solutions. This technology enables small enterprises to initiate operations with limited capital and expand in accordance with their evolving requirements. Although cloud computing is still nascent, it possesses the capacity to fundamentally revolutionize the IT industry. It encompasses essential topics including architecture, implementation strategies, foundations, and potential future research domains. Researchers are endeavoring to resolve design challenges and introduce innovative approaches to this dynamic field. As cloud computing evolves, it is expected to remain a catalyst for technical progress.

Keywords: *Cloud computing, Data centers, Virtualization.*

Date of Submission: 27-06-2025

Date of Acceptance: 10-07-2025

1. INTRODUCTION

The proliferation of the Internet and advancements in processing and storage have made formerly expensive and scarce computing resources more accessible and powerful than ever before. The ability to rent and return servers and data centers online is what makes cloud computing viable. Customer management of cloud platforms and resource distribution is made possible through leases from service providers. Organizations like Google, Amazon, and Microsoft were the first to offer cloud computing. The benefits of cloud computing for companies are detailed below.

Use is the primary determinant of pricing for cloud services. Using the cloud, service providers can avoid the need for servers. This cloud leases out computational resources.

Cloud computing is economical because it quickly assigns and releases computer

resources. This meant that service providers could do without peak capacity. Reduced demand for services may allow for the reallocation of resources to more cost-effective areas. Companies specializing in infrastructure offer scalable data centers. As an illustration of how rapidly service providers can expand to fulfill demand spikes, the "flash crowd" has emerged. A few words require further punctuation. You can access the most popular cloud services online. They are present in a wide variety of comparable devices. Electronic devices can be programmed. When companies outsource their cloud-based service infrastructure, they cut down on maintenance expenses and the likelihood of hardware failure. Training employees and keeping gear up to date will cost organizations less.

Cloud computing has many positive effects on the IT business, but it has also introduced some new challenges. We take



a look back at the cloud computing architecture, its applications, and the holes in our present knowledge. The challenges of cloud computing design will be examined.

Give some ideas on where to go from here for additional studies in this intriguing field.

2. RELATED TECHNOLOGIES

These technologies are comparable to cloud computing in that they employ networks within the computer model. Grid computing was developed in response to the high processing power demands of scientific computing. In cloud and grid computing, tasks are executed by dispersed servers. Virtualization is the key to cloud computing, which simplifies resource sharing and dynamic provisioning.

Utility computing provides resources on an as-needed basis. As an example of utility computing, consider the cloud. maximizing output by implementing pricing based on the cost of services. By adopting utility-based pricing and on-demand provisioning, service providers can optimize resource consumption while cutting costs.

Virtual machines can obfuscate or expose hardware. simulation of essential software applications. Computers are necessary for virtual reality servers. Virtualization allows cloud computing to dynamically allocate resources from the device pool to applications.

By creating devices that react to both internal and external inputs, IBM was a trailblazer in autonomic computing in 2001. Automation software simplifies the management of computer resources. Cloud computing saves money by providing

resources independently, rather than by simplifying infrastructure.

The resources of cloud computing are available online and whenever needed. There are three different types of computing: parallel, autonomous, and grid. Completing its prerequisites comes with its own set of pros and cons.

3. CLOUD COMPUTING ARCHITECTURE

Cloud computing's architectural, operational, and commercial components are examined.

A layered model of cloud computing

The diagram shows the four levels of cloud computing: hardware/datacenter, infrastructure, platform, and applications. Just one. To make things easier for you, the breakdowns are displayed below.

The hardware department oversees the electrical infrastructure, networks, servers, and environmental systems. Machines are the building blocks of data centers. Switches, routers, and fabric link hundreds of servers housed on racks in data centers. Hardware includes features such as configuration, fault tolerance, energy consumption, thermal regulation, and traffic control.

The virtualization layer integrates storage and processing power from several physical servers using technologies like Xen, KVM, and VMware. Virtualization makes the infrastructure layer vital for cloud computing capabilities like dynamic resource allocation.

At its most fundamental level, an infrastructure consists of an operating system and application frameworks. Virtual machines are able to access programs that are located at the platform layer. Databases, business logic, and

storage can all be accessed using the platform-layer APIs provided by Google App Engine.

At the very top of the cloud, in what is called the "application layer," are all the applications. When apps in the cloud can scale on their own, they improve performance, availability, and cost-effectiveness. Contrast modular cloud computing with server farms. Due to the weak interlayer connections, each layer evolves separately. You can see this in action with the Open Systems Interconnection (OSI) suite of network protocols.




Fig. 1 Infrastructure that is hosted in the cloud..

The scalability and affordability of cloud computing make it an attractive option for many different types of applications.

Business model

Computing in the cloud. This allowed us to get the required software and hardware. The architecture builds upon itself with each successive layer. Intersecting layers are consumed. Three acronyms that sum up cloud services: SaaS, PaaS, and IaaS. Internet as a service makes use of virtual machines. Cloud computing companies offer their customers infrastructure as a

service. Amazon Elastic Compute Cloud (EC2), GoGrid, and Flexiscale are few of the IaaS providers.

Developers get access to the PaaS OS and framework. One example of a platform as a service is Google App Engine. Another is Microsoft Azure.

Application service provided by the cloud on demand. Companies like Rackspace, SAP Business By Design, and Salesforce.com provide software as a service.

Figure. This picture is a representation of cloud computing. Two. With the use of the cloud computing ecosystem's tiers, PaaS providers can access IaaS clouds. There is a vast array of cloud computing services offered by Google and Salesforce. The suppliers of PaaS and IaaS are referred to as "cloud" and "infrastructure" in this context.




Fig. 2 A cloud-based business model
TYPES OF CLOUDS

Moving a business app to the cloud could be challenging. Dependability and safety are not prioritized by all retailers. There are benefits and drawbacks to cloud computing.



Providers can pool their resources by utilizing open networks. Public clouds help providers save money and lessen their risk. Public cloud services still don't meet the needs of many businesses because of problems with data, network, and security management.

In a private cloud, only one company uses it. Companies and organizations offering their services have the option to construct their own private clouds. A private cloud enhances and secures data processing. Cloud computing has come under fire for its parallels to proprietary data centers and its low entry criterion.

Hybrid clouds use the best features of both public and private clouds to overcome their limitations. Hybrid clouds combine the best features of public and private cloud computing. Both public and private clouds are inflexible, unlike hybrid clouds. Private clouds provide for the safe and scalable implementation of on-demand data and application services. The development of hybrid clouds is hindered by the need to differentiate between private and public clouds.

There are several benefits to using a VPC instead of a conventional cloud installation. "Public clouds" refer to both physical and virtual private clouds. By utilizing VPNs, service providers have the ability to personalize firewall and VPC configurations. One way to improve the VPC is by virtualizing the server networks, software, and hardware. The virtualized network layer of a virtual private cloud (VPC) facilitates the migration of several organizations to cloud computing.

Most providers of services favor cloud computing. Scientific computational applications can be implemented cheaply using the public cloud. Unpredictability

characterizes cloud formation and dissipation.

Hybrid technologies may become increasingly common in the workplaces of the future. The use of VPCs, or virtual private clouds, has been on the rise since 2009.

4. CLOUD COMPUTING CHARACTERISTICS

Cloud computing offers several advantages over service computing.

Vendor services for cloud data centers are a collaborative effort between customers. Companies offering infrastructure and services often struggle with concerns of efficiency and management. The proprietors of the cloud layers can concentrate on their own objectives when they've been given responsibility. Multitenancy reduces stakeholder participation.

Distributing computing resources to consumers dynamically is made possible by the infrastructure provider through shared resource pooling. With dynamic source assignment, infrastructure service providers can keep better tabs on resource expenditures. Virtual machine migration allows infrastructure as a service companies to save money on cooling and energy costs while sharing servers and optimizing resources.

The Internet has made cloud services accessible from any location. You can access cloud services from any internet-connected device, including smartphones, PDAs, and desktop computers. In an effort to localize their networks and increase efficiency, several cloud providers have set up data centers all around the globe. When there are many of easily accessible locations, service providers may be able to

work more efficiently.

In theory, cloud computing might pave the way for a more service-oriented society. Essential is the administration of services. Service Level Agreements govern cloud computing services like Infrastructure as a Service (IaaS), Platform as a Service (PaaS), but also Software as a Service (SaaS). Service Level Agreements (SLAs) are a must for every service provider.

Cloud computing provides assistance whenever needed. Dynamic resource provisioning allows service providers to save money by purchasing resources only when they are needed. With dynamic resource allocation and release, service providers gain greater agency over resource management. The "flash crowd" is a target market for several companies that deploy administrative automation software.

One benefit of using cloud services is the "pay as you go" model. A wide range of services have different prices. Virtual servers are rented by IaaS providers to SaaS enterprises by the hour. Salesforce's on-demand CRM may have per-user costs. A cost-cutting measure, the utility-based pricing model charges consumers only for the resources they really use. It is becoming increasingly difficult to monitor spending. Through VKernel, cloud clients are able to track resource consumption, assess expenses, and implement modifications.

5. STATE-OF-THE-ART

Implementing cutting-edge cloud applications. Cloud computing is a relatively new idea. After that, we'll examine data and services hosted on the cloud.

CLOUD COMPUTING TECHNOLOGIES

This section discusses cloud computing.

a. Architectural design of data centers

In order to store and handle data, cloud computing makes use of servers, switches, and routers housed in expansive data centers. Distributed computing is made more efficient using deep network design. Contemplate long-term viability and adaptability. Tier networks have been researched in big data centers. Figure. This graphic represents the data center's processing, storage, and networking components. The number 3. Through access layers, server rooms are linked to the internet. Each building houses twenty to forty servers that are connected via 1 Gbps access switches. There is redundancy in the connections between the access switches and the aggregation switches.




Fig. 3 Putting Together a Data Center Network

A speed of 10 Gbps is within reach. The aggregation layer is responsible for regulating the workload distribution



between domains, locations, and servers. A large number of aggregation switches make up a routed network with core-layer failure tolerance. Core routers manage all traffic originating from and going out of the data center.

Designers of networks often make use of inexpensive Ethernet switches and routers. No business is too big or too little to implement a multilayered network. To achieve the objectives of data center network design, server and communication allocation should be done apart from network design concerns. A fast network card is essential for data center hosts to communicate with one another.

Virtualization allows for the easy transfer of virtual machines between different networks. Companies that provide cloud hosting have the ability to manage the temperature and electricity usage of data centers, change connection settings for hosts that are heavily networked, and relocate virtual computers in order to use statistical multiplexing. In order to make the migration of virtual machines easier, optimize your network infrastructure.

Massive setbacks are inevitable. A well-planned network should be able to withstand failures in servers, connections, and racks. There must be robust physical connection for both multicast and unicast transmissions. The growing number of users and servers necessitates the growth of the network.

It is vital that Ethernet and IP networks are compatible. The new layout must be compatible with most data center Ethernet and IP devices.

Another innovation in technology is the introduction of data centers that are modular and housed in freight containers. To link thousands of servers together,

MDC networks use switches. Geodispersed MDCs near to major population centers are enhanced using multipurpose software. Redundancy is provided by MDC since blackouts, earthquakes, and instability do not occur frequently. Guo and colleagues favored server-centric recursive multi-dimensional codes.

b. Distributed file system over clouds

To provide fast and reliable data access, Google's File System (GFS) makes use of massive clusters of inexpensive general-purpose computers. Most files are viewed in 64 MB chunks and rarely changed. Throughput, latency, and server stability are all enhanced in a data center environment with GFS.

One way in which HDFS enhances GFS is by facilitating the storing of big files across multiple systems. For safety reasons, data replication between servers is certainly an option. Nodes in the network, including GFS, store data. Users all around the network can more easily get data blocks from HDFS nodes thanks to the block protocol. Any and every data sent via HTTP is captured by the client or browser. Communication between nodes facilitates data sharing, archiving, and replication.

c. Distributed application framework over clouds

Visual effects and research on economic trends are two uses for server clusters in today's data centers. Web applications that use HTTP utilize Java EE.

To manage the distributed processing of big data clusters, Google created MapReduce. Clients use the master to receive MapReduce jobs. Smartly distributed across the data center, task nodes are assigned work by the Master.



All information and host nodes are available to the master. When one data node goes down, the other rack nodes step in to fill the void. Decreases in core backbone traffic lead to gains in network throughput. postponing responsibility in an effort to evade it. With the Master down, the show must end. Master processes rely on file system logs. It seeks out these data at startup. Hadoop MapReduce owes a debt to Google. Organizations can utilize Hadoop MapReduce to analyze massive volumes of data.

COMMERCIAL PRODUCTS

There is a comparison of the best cloud computing companies.

a. Amazon EC2

The processing, storage, and other features offered by AWS allow consumers and businesses to build apps and services rapidly and at a reasonable cost. On AWS, you can use REST, SOAP, or HTTP.

By utilizing APIs and other technologies, cloud clients are able to deploy and control instances of Amazon Elastic Compute Cloud (EC2) servers. Elastic Compute Cloud uses Xen. New or revised code can be delivered by users using instances. Significant modifications can be seen in recent machine images. Feel free to make copies anytime you like. Running programs on hardware, such as EC2 instances, puts nearly all control over software design in the hands of the user. This update complicates Amazon's automated resource scaling.

Split up the EC2 instances. There are availability zones and regions that make up EC2. Territory refers to areas of designated availability. Isolating availability zones within the same region reduces latency between them.

S3 is the way to go for EC2 machine photographs. S3 "buckets" may hold data ranging from one byte all the way up to five terabytes. You can get a lot done with just one bucket. Find the perfect spot while keeping user delay and expense in mind.

Integrating AWS with existing IT infrastructure is a breeze with Amazon Virtual Private Cloud. Businesses can utilize AWS computer resources including firewalls, intrusion detection systems, and network security systems through Amazon Virtual Private Cloud (VPC).

In order to give cloud customers with insights in near real-time, Amazon CloudWatch pulls raw data from AWS partner services like Amazon EC2. EC2 keeps tabs on your CPU usage, disk I/O, and network tardiness.

b. Microsoft Windows Azure platform

Three parts make up Microsoft's Azure cloud platform. While SQL Azure is built on top of Microsoft SQL Server, .NET Services provides a distributed framework for development in the cloud and on-premises.

Software hosted on-premises and in the cloud are both supported by Windows Azure.

Create Windows Azure applications. Microsoft Windows is compatible with Visual Basic, C, C#, and the .NET Framework. The capabilities of Windows Azure applications are constantly improving. Web applications can be built in Lulls using either ASP.NET or WCF, or even a mix of the two. Microsoft Azure queues, tables, and blobs can be accessed over RESTful HTTP/HTTPS.

SQL Azure hosts the database and service known as "Huron" Data Sync. SQL Azure's backend database is SQL Server.



You can use ADO.NET or another Windows data API to get access. You can access the cloud using local programs. Relational data synchronization with the on-premises system "Huron."

Table 1 Popular things are compared

Table 1 A comparison of representative commercial products

Cloud Provider	Amazon EC2	Windows Azure	Google App Engine
Classes of Utility Computing	Infrastructure service	Platform service	Platform service
Target Applications	General-purpose applications	General-purpose Windows applications	Traditional web applications with supported framework
Computation	OS Level on a Xen Virtual Machine	Microsoft Common Language Runtime (CLR) VM; Predefined roles of app instances	Predefined web application frameworks
Storage	Elastic Block Store; Amazon Simple Storage Service (S3); Amazon SimpleDB	Azure storage service and SQL Data Services	BigTable and MegaStore
Auto Scaling	Automatically changing the number of instances based on parameters that users specify	Automatic scaling based on application roles and a configuration file specified by users	Automatic Scaling which is transparent to users

Computation

It is possible to distribute software using .NET Services. Cloud services and businesses can utilize Access Control to incorporate user authentication. Web services for Service Bus can be accessed using both on-premises and cloud-based applications. To reach endpoints, clients use URIs.

Software, virtual machines, and servers in data centers are all managed by fabric controller software. We have sent the software's XML configuration parameters. By utilizing this data, fabric controllers can allocate servers to new applications.

c. Google App Engine

Cloud application deployment is made easy with Google App Engine. We support both Python and Java. App Engine is compatible with Google's own web framework, as well as web2py, Django, CherryPy, and Pylons. Google is responsible for cluster setup, monitoring, failover, and application activation. Data storage and retrieval in Big Table, as well

as HTTP requests, are managed by the APIs. App Engine developers have read-only access to the file system.

Compare and contrast the three main cloud providers in Table 1 based on their autonomous scaling, processing power, storage capacity, application categories, and utility computing. There are a variety of data encapsulation and management techniques used by cloud services. Several cloud options may meet the needs of the firm.

6. RESEARCH CHALLENGES

The concept of cloud computing is still in its early stages, but it is quickly becoming popular. We don't yet know what these applications mean for industry. Cloud computing research in the future will face challenges.

Automated service provisioning

The flexibility to gain or lose access to resources on demand is a major perk of cloud computing. The reallocation of cloud resources is done by service providers in order to achieve SLOs while reducing costs. A variety of strategies are available to service providers. There is a clear correlation between physical assets and service excellence. Demand fluctuations, or the "flash crowd effect," can cripple internet services.

Using machines to help out is nothing new. The idea of dynamic resource allocation is not new to web applications. An application performance model determines the number of application instances required to satisfy QoS standards at each tier, automated resource allocation, and periodic demand and resource projections are typical components. Through the integration of statistical machine learning with queueing and



control theory, the application's effectiveness is modeled.

There are two main approaches to managing resources: proactive and reactive. Prior to utilizing assets, they must be deployed. Unpredictable shifts in demand are met with a reactive approach. In today's fast-paced business world, both approaches to managing resources are crucial.

Virtual machine migration

By dispersing virtual machines across a data center, cloud computing evenly distributes processing demands. Data center provisioning that is both fast and reliable allows for the migration of virtual machines. Change to an off-site PC that is currently running. Transferring "live" VMs from Xen to VMware took a few tens of milliseconds not long ago. Live virtual machine migration, according to Clark et al., reduces process migration problems.

Server consolidation

This is not a complete or accurate statement. Hotspots are alleviated by virtual machine migration, however there are substantial hurdles. For the rapid detection and relocation of workload hotspots, existing approaches fall short. Considering the physical server's and apps' resources, memory allocation must be secure while still being efficient. Streamlining operations and reducing power consumption are two benefits of moving server resources to the cloud. Virtual machines (VMs) can be migrated live from idle servers using VM consolidation technologies. Some have drawn parallels between data center server consolidation and vector bin packing, an NP-hard problem. The problem was fixed using heuristics.

Energy management

The importance of communication has just recently been acknowledged. Server consolidation must not compromise application performance. Virtual platforms require specific resources. Excessive consumption of bandwidth, memory cache, or disk I/O by virtual machines during server consolidation may lead to resource congestion. Monitoring alterations in the VM footprint may be essential for server consolidation. Ultimately, enhance the efficiency of your constrained assets.

Another target is the energy efficiency of cloud computing. Electricity and refrigeration account for about fifty percent of a data center's running costs. By 2020, data centers were projected to consume 18% of the electricity utilized in the United States, a substantial rise from the 1.5% consumed in 2006. Energy efficiency is a vital priority for infrastructure companies. Data centers can reduce their power use for legal and environmental purposes.

Eco-friendly data centers are becoming progressively prevalent. There are various approaches to resolve this problem. CPUs with lower power consumption may suffer a decline in performance or fail to operate entirely. Server consolidation and advanced workload scheduling are employed to save energy by deactivating workstations. Recent studies have concentrated on the advancement of energy-efficient network protocols and architectures. Each of these solutions reduces the application's functionality to conserve electricity. Insufficient research has been noted concerning cloud performance and effectiveness.



Traffic management and analysis

Traffic analysis plays a crucial role in data centers. Many websites employ traffic statistics to enhance the user experience. In order to plan and manage networks, operators must understand traffic patterns. Technologies that measure and analyze data center traffic provide a variety of obstacles for companies and ISPs. To start, the link density is far higher than in ISP and corporate networks, rendering the aforementioned techniques useless. 2) In a modular data center with thousands of servers, the majority of algorithms are only capable of producing traffic matrices for a small number of hosts. Road map The company's network and Internet traffic are affected by the reduction in workload. Programs are associated with storage, processing power, and network utilization. Data centers seldom keep tabs on network activity. Data center traffic parameters, like concurrent flows and flow volumes, are utilized by Greenberg and associates to construct networks. Using SNMP traces from routers, Benson et al. examine traffic on the edge of data centers.

Data security

Protecting sensitive research data is an absolute must while using cloud computing. Because service providers lack direct access to data center security, they depend on infrastructure vendors. When it comes to remote security, virtual private cloud providers can only set rules; they can't enforce them. To guarantee the security of the application, the infrastructure provider must encrypt data access and transfer and conduct audits. Remote attestation obfuscates cryptographic keys.

The Trusted Platform Module (TPM) uses

its private key to encrypt the system state and create an irrefutable system summary during remote attestation. Because cloud-based virtual machines are inherently temporary, remote attestation presents challenges. Every layer of the cloud architecture relies on trust mechanisms. Hardware-based security is made easier with TPM. Users of virtual reality systems need to have faith in the technology. Rest assured, only reliable servers should be used for virtual machine transfer. Building trust has been the subject of recent studies.

Software frameworks

Applications that rely heavily on data run smoothly in the cloud. Hadoop and other MapReduce frameworks make this software scalable and fault-tolerant. The efficiency of MapReduce jobs and the resources needed to complete them vary across different types of applications, according to a recent research. When it comes to Hadoop, sorting is mostly dependent on I/O and grep is mostly dependent on CPU. On Hadoop hosts, the number of virtual machines can be limitless. A server's bandwidth is limited by the number of virtual machines it hosts. In order to make MapReduce applications more efficient and cost-effective, changes to the setup and scheduling are required. Application execution is accelerated by removing bottlenecks. The most difficult parts of working with Hadoop are dynamic scheduling and modeling the performance of both online and offline tasks.

Using MapReduce frameworks that are energy efficient is one possibility. Until they get fresh tasks, idle Hadoop nodes do nothing. It's crucial to keep energy consumption low for HDFS and Hadoop. There might be a conflict between



performance and energy efficiency. The best outcome trade-off can only be determined with more research.

Storage technologies and data management

The method makes use of GFS and HDFS, two large-scale file systems accessible over the internet. These file systems differ from dispersed file systems in terms of storage, accessibility, and application programming interface. When it comes to older programs and file systems, non-POSIX interfaces just won't cut it. Lots of studies have looked into this. The MapReduce framework and the primitives for scalable, concurrent data access in IBM's GPFS cluster file systems were created by Patil and colleagues.

Novel cloud architectures

Data centers are the physical locations of commercial clouds. Even though they are designed for administration, data centers are expensive and energy-intensive. Smaller data centers require less expensive and more efficient cooling systems, but they are more distributed, cheaper to build, and use less power, according to a recent research. Geographic diversity is advantageous for real-time applications like gaming and content dissemination. The capacity of nano-data centers to transmit video was examined by Valancius et al.

An unrelated topic is the hosting of cloud applications on infrastructure that the user provides. An increase in output with a decrease in expenditure is possible with dedicated servers or cloud systems run by volunteers. While planning, think about things like resource management and staff turnover. The reasons behind these

classifications have been studied by researchers.

7. CONCLUSION

The cloud keeps your systems and networks up to date. With cloud computing, utility computing is made possible. Regardless of its benefits, cloud-based technologies are always evolving. This research mainly focuses on three main areas: autonomous resource provisioning, power management, and security. We are confident that this research will advance the field.

The research examined cloud computing's foundations, architectural designs, key characteristics, crucial technologies, and research objectives. We expect that as cloud computing becomes more popular, this work will shed light on design considerations and encourage more research along these lines.

REFERENCE

1. Al-Fares M et al (2008) A scalable, commodity data center network architecture. In: Proc SIGCOMM
2. Amazon Elastic Computing Cloud, aws.amazon.com/ec2
3. Amazon Web Services, aws.amazon.com
4. Ananthanarayanan R, Gupta K et al (2009) Cloud analytics: do we really need to reinvent the storage stack? In: Proc of HotCloud
5. Armbrust M et al (2009) Above the clouds: a Berkeley view of cloud computing. UC Berkeley Technical Report
6. Berners-Lee T, Fielding R, Masinter L (2005) RFC 3986: uniform resource identifier (URI): generic syntax, January 2005
7. Bodik P et al (2009) Statistical



-
- machine learning makes automatic control practical for Internet datacenters. In: Proc HotCloud
8. Brooks D et al (2000) Power-aware microarchitecture: design and modeling challenges for the next-generation microprocessors, IEEE Micro
 9. Chandra A et al (2009) Nebulas: using distributed voluntary re- sources to build clouds. In: Proc of HotCloud
 10. Chang F, Dean J et al (2006) Bigtable: a distributed storage system for structured data. In: Proc of OSDI