2.3 Performance Analysis of Cloud Computing Architectures Using Discrete Event Simulation

Performance Analysis of Cloud Computing Architectures Using Discrete Event Simulation

John C Stocker & Andrew M. Golomb
Booz Allen Hamilton, Inc.
stocker_john@bah.com golomb_andrew@bah.com

Abstract. Cloud computing offers the economic benefit of on-demand resource allocation to meet changing enterprise computing needs. However, the flexibility of cloud computing is disadvantaged when compared to traditional hosting in providing predictable application and service performance. Cloud computing relies on resource scheduling in a virtualized network-centric server environment, which makes static performance analysis infeasible. We developed a discrete event simulation model to evaluate the overall effectiveness of organizations in executing their workflow in traditional and cloud computing architectures. The two part model framework characterizes both the demand using a probability distribution for each type of service request as well as enterprise computing resource constraints. Our simulations provide quantitative analysis to design and provision computing architectures that maximize overall mission effectiveness. We share our analysis of key resource constraints in cloud computing architectures and findings on the appropriateness of cloud computing in various applications.

1.0 INTRODUCTION

1.1 Motivation

Organizations migrating to cloud computing are faced with the challenge of either negotiating service level agreements with a cloud provider or allocating limited resources to develop their own cloud computing implementation. In weighing appropriate cloud providers or architectures in developing their own implementations, organizations should take care to maximize their overall mission performance while minimizing cost.

Government organizations are motivated by the cost savings that can be realized through a cloud computing, but are generally hesitant to use established commercial services because of concerns over privacy and information security, [2], [4], [9]. As a result of these issues, real and perceived, many government clients have focused on private and community cloud architectures. Determining the appropriate size for a private or communication cloud is a novel problem in the cloud computing community, as commercial sector enterprises growing more confident in leveraging commodity cloud services from established vendors, [2]. Details on the capacity and infrastructure among commercial cloud service providers are closely held as trade secrets, and cannot benefit organizations looking to develop their own cloud infrastructure. Some providers even protect the number of data centers they operate, [5].

Appropriately resourcing cloud infrastructure that will both enable rapid elasticity as needed to meet user demand while not over building is a significant challenge. We present a quantitative cloud computing architecture effectiveness and performance model framework to support the design and analysis of cloud architecture implementations. Dynamic modeling is necessary to ensure that enterprise effectiveness is maintained as mission requirements change over time. We anticipate two primary applications of our model framework:

- *Cloud infrastructure design*: Enable quantitative requirements analysis for the specification of cloud characteristics and infrastructure design.
- *Service Level Agreement specification*: Model anticipated service usage across the organization determine required quality of service (QoS) for cloud services.

1.2 Analyzing Cloud Effectiveness and Performance

We developed a cloud computing effectiveness and performance model framework to analyze the effectiveness of various cloud computing architectures in meeting enterprise computing requirements by applying operations research techniques. The analytic model coupled with a companion cost analysis model [1] will enable the development of comprehensive migration strategies that right-size investments in developing private, community, and public cloud computing solutions. While previous research has analyzed the performance of cloud computing [10], that work focused on cloud service requirements, and not detailed analysis of various cloud architectures. To the best of our knowledge, we are the first to develop a generalized approach that can provide analysis of detailed cloud infrastructure characteristics.

2.0 EFFECTIVENESS AND PERFORMANCE MODEL

The effectiveness and performance model combines a dynamic workforce business process model with a cloud computing architecture resource model. Figure 1 demonstrates a top level view of model components. The model framework can be configured to simulate specific processes and a variety of cloud environments. Discrete event simulation is used to execute business processes in the context of cloud computing architectures. The model is based on a discrete event simulation model built using ExtendSim® [3] developed by Imagine That Inc.

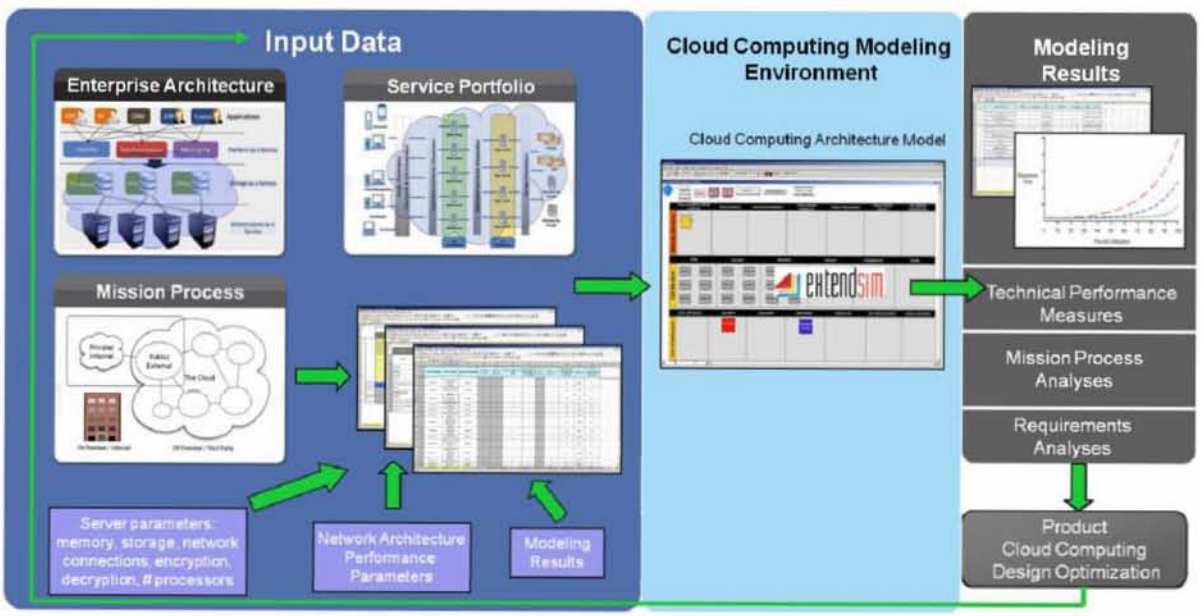

Figure 1. Cloud Computing Modeling Methodology

2.1 Business Process Model

The business process model generates compute service requests that are dispatched to the cloud for execution. The process model considers the number of users, their roles, and the distribution of service requests based on their roles. Services are configured to reflect the organization's mission and can include generic and specific activities such as: email, instant messaging, video chats, web browsing, document review and composition, budget and financial analysis, and scientific computing.

The business process model is currently comprised of two top level models:

- Office
- Service

Numbers of staff are affiliated within offices to model the overall organization effectiveness in community cloud architectures. Network characteristics between users and cloud computing resources are specified at the office level to accurately model differences in network connectivity between users of community clouds. Table 1 describes the characteristics of the office model:

Office site name
Number of staff members
Communication link characteristics - Maximum uplink bandwidth (Mbps) - Maximum downlink bandwidth (Mbps)

Table 1. Office model

The business process model has sophisticated capability to generate service demand based on the rhythm in the workforce and enterprise needs. The process model can also represent surges in service demand that may be triggered by mission events. This capability is crucial in analyzing the benefits of cloud elasticity in meeting mission requirements. Table 2 defines the characteristics of the software or service modeled.

Application or service name
Request rate
Data transferred (KB)
Computational requirements (MIPS)
QoS requirements - Allowable response latency

Table 2. Service model

The service model characterizes request rate in terms of on demand, hourly, daily, weekly, monthly, and quarterly. On demand requests are specified using a statistical distribution of demand rate or interarrival time. Data transferred and computational requirements are also statistical distributions that define average and peak values.

Using statistical distributions permits significant flexibility in modeling various services and processes. Distributions can be varied (e.g., Gaussian, Poisson, triangular) to most accurately represent the specific enterprise and mission being modeled. The business process model framework is flexible and can represent diverse work streams throughout the enterprise.

2.2 Cloud Computing Architecture Model

The cloud computing architecture model captures overall cloud configuration and resources. The model has a flexible design that facilitates simulation of various cloud computing architecture deployment models identified by the National Institute of Standards and Technology [6], also shown in Figure 2 [7].

- Community
- Hybrid
- Private
- Public

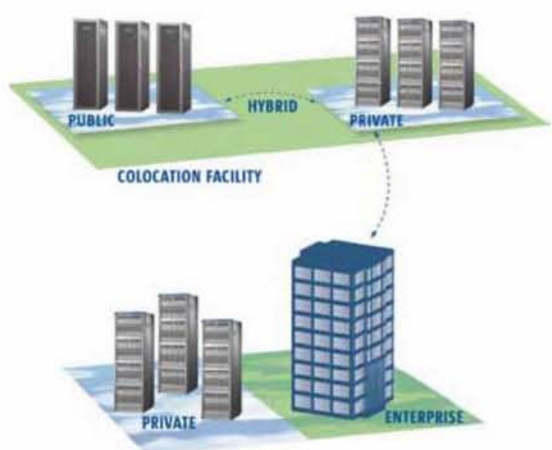

Figure 2. Cloud Computing Architectures

The cloud computing architecture model represents the underlying cloud infrastructure, including network and servers. The model assumes uniform performance with respect to the three cloud service models: Software as a Service (SaaS), Platform as a Service (Paas), and Infrastructure as a Service (IaaS). The distinctions between these service models are largely based on revenue models and not technical distinctions that affect overall performance of a cloud-based service solution.

NIST further describes the characteristics of cloud computing environments, including:

- On-demand self-service
- Broad network access
- Resource pooling
- Rapid elasticity
- Measured service

The scope of model focuses on providing an analysis of necessary cloud computing node resource characteristics, such as network access, processing capacity, memory, and storage. The model assumes that the nodes are homogenous. The internal network resources between cloud nodes are also modeled and can represent various server interconnectivity solutions. Table 3 specifies the characteristics captured in the cloud architecture model.

Number of Nodes
Node processing capacity
Node memory
Node storage capacity
Node storage access speed
Internode network characteristics

Table 3. Cloud model

A limitation of the cloud architecture model in its current state is its inability to represent hybrid cloud solutions that incorporate multiple instances of cloud environments into a single implementation. This limitation restricts current analysis to consider only private, community, and public infrastructure. Future work is planned to incorporate the ability to represent hybrid implementations that consider combinations of private, community, and public cloud environments.

It is important to note that the cloud architecture model framework is also flexible enough to represent traditional fixed hosting infrastructure to permit comparisons of service performance between cloud and traditional hosting environments.

3.0 EFFECTIVENESS AND PERFORMANCE METRICS

3.1 Measures of Effectiveness

Measures of effectiveness focus on the ability of the cloud infrastructure to meet service demand in a timely fashion and enable the organization to perform its mission.

3.1.1 On-Time Response Rate

On-time response rate is the cumulative distribution of service requests that are completed and responded to within the required latency. Acceptable latency is service and mission specific and typically varies in magnitude from fractions of seconds to minutes.

3.1.2 Service Latency

Service latency is the delay experienced by users in processing service demand. This measure is an indicator of services along the critical path in mission execution and an indicator of potential idleness in the workforce.

3.2 Measures of Performance

Measures of performance characterize cloud resource utilization in meeting the service demand. The measures of performance gauge cloud elasticity in the context of an enterprise's mission requirements.

3.2.1 Cloud Processing Utilization

Cloud processing utilization, as measured by the usage of overall cloud processing capacity over time and across all services in the mission workflow, is an indication of overall processing demand.

3.2.2 Cloud Storage Utilization

Cloud storage utilization, as measured by the usage of cloud storage capacity over time and across the workflow, is an indication of overall storage demand.

3.2.3 Cloud Communications Utilization

Cloud communication utilization is the percentage of communications resources over time. Individual measures are recorded for each office communications link as well as the cloud communications infrastructure.

4.0 ANALYSIS

Quantitative cloud architecture analysis using the effectiveness and performance model will help answer fundamental design questions and support trade studies that refine specific cloud technologies and solution sets. Analyzing the impact of cloud architectures characteristics on the enterprise is crucial to designing specifying acceptable cloud architecture implementations.

The cloud computing effectiveness and performance model will help answer design questions such as:

- Is a cloud architecture appropriate for the mission?
- Is a private, community, or public cloud best suited to meet performance requirements?
- Does the cloud have sufficient elasticity for all enterprises, offices, users, and missions?
- What is the right size cloud to meet mission requirements and minimize cost?
- Do communications resources need to be upgraded when services are migrated to a cloud architecture?

These questions are best addressed by considering the overall on-time response rate across all service requests among all users relative to the performance metrics. Ideally, the on-time response rate is 100%, however, given the scale of requests across the enterprise and resulting peaks of activity it's impractical to build a cloud that can scale rapidly enough to meet 100% on-time response rate at all times.

4.1 Example sensitivity analysis

We conducted sensitivity analysis of on-time response rate based on some of the key cloud characteristics for a sample private cloud infrastructure. The sensitivity analysis of on-time response rate was performed by varying cloud computing architecture resource parameters. We looked for deflection points where there were diminishing returns in the overall on-time response rate relative to additional cloud resources: either more additional nodes or increased bandwidth.

	Daily usage per user		Bandwidth Sent/Received (KB)		Allowable Latency (seconds)	Million CPU Instructions Required	
Service Request	Average	Maximum	Average	Maximum		Average	Maximum
Email	30	500	4	1000	2	0.1	10
Database access	3	5	3000	60000	120	1000	10000
Instant messaging	50	500	1	2	2	0.05	1
Email w/ attachment	5	20	2000	20000	60	1	10
Wiki/Blog entry	1	4	2	40	5	1	10
Office app usage	30	100	8	500	1	10	100
Model / simulation	8	16	200	1000000	2400	100000	10000000
Financial Report Gen	1	3	10	40	1200	10000	100000
Video chat	5	20	6	40	1	1000	10000
Collaboration site	10	20	20	100000	5	10	100

Table 4. Services modeled in sample analysis

Our example sensitivity analysis considers a single enterprise office with 5,000 staff members using a private cloud infrastructure. In this sample sensitivity analysis we assumed that all users had uniform workflow and used services in equal distributions. It is reasonable to approximate usage at the service level across the enterprise instead of the user role level. Table 4 provides a summary ofthe service characteristics modeled in this example analysis.

Table 5 describes the cloud architecture characteristics used in this example analysis. Given the private cloud studied in this example, the bandwidth is the data rate between the office and the cloud implementation.

Bandwidth (Mbps)	5,000
Number of nodes	500
CPU Processor speed (GHz)	2
Instructions per cycle	16
Storage capacity per node (TB)	0.1
RAM (GB)	1

Table 5. Cloud architecture modeling in sample analysis

In our sensitivity analysis we independently varied the available bandwidth between the office and cloud and the number of nodes to assess the impact of each on the on-time response rate. Other variables were held at values that would not impact effectiveness. In the sensitivity analysis, the cloud architecture was fixed at 500 nodes.

4.2 Example sensitivity analysis

We observed a high sensitivity on bandwidth in our example analysis. Figure 3 shows the sensitivity of the on-time response rate on bandwidth.

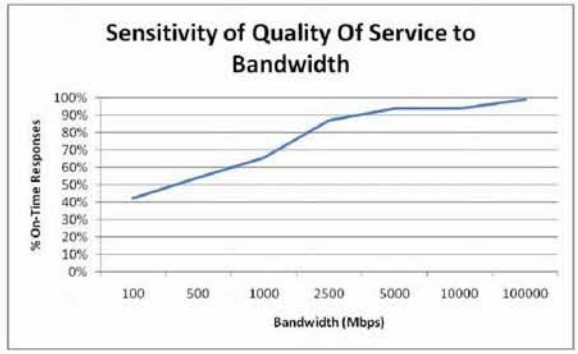

Figure 3. Bandwidth sensitivity sample analysis

The on-time response rate shown varies from approximately 42% to 94% as bandwidth is increased from 100 Mbps to 5,000 Mbps. These results are not surprising. The services modeled were indeed activities that required a significant amount of data to be transferred between users and the cloud. Figure 4 provides additional insight on the response rate for categories of services.

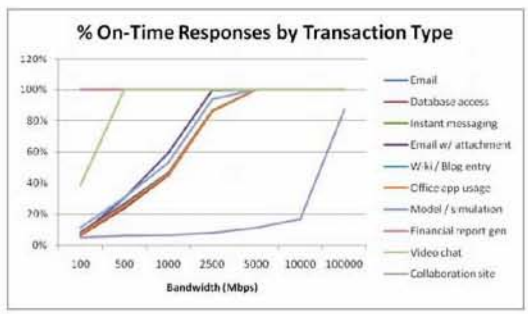

Figure 4. Bandwidth sensitivity by service

5.0 CONCLUSION AND FUTURE WORK

We have developed an effectiveness and performance model framework to assess the impact of various cloud infrastructure characteristics on overall enterprise effectiveness.

Our preliminary analysis indicates that cloud architectures may require the investment of network infrastructure to provide high quality of service to applications where a significant amount of data must be transferred.

Additional effectiveness considerations to model in the future include assessing cloud availability and analyzing various cost options in maximizing availability through large-scale implementations of commodity hardware versus more deploying limited quantities of more expensive hardware.

In the future we plan to add to the fidelity of the business process model by considering the categorization of staff members by role and modeling service request generation on a role-by-role basis. We also plan to incorporate a hierarchical business process model in which required capabilities can be decomposed into dependent software applications and services. We will leverage existing work in modeling business processes using industry standard representations such as the Federal Enterprise Architecture Framework (FEAF), [8].

The cloud computing effective and performance model provides a quantitative approach to right-size private and community cloud investments. Coupling the cloud effectiveness model with the existing Booz Allen cost model [1] will provide an overarching analytical approach to shape cloud migration strategies. This integrated analytic capability would provide detailed performance and cost analysis of an organization's mission operations in a variety of cloud environments.

6.0 REFERENCES

[1] T. Alford, G. Morton, "The Economics of Cloud Computing: Addressing the Benefits of Infrastructure in the Cloud," Booz Allen Hamilton, McLean, Virginia, 2009. [Online]. Available: http://www.boozallen.com/media/file/Economics-of-Cloud-Computing.pdf

[2] Gartner Research, "Gartner Says Worldwide Cloud Services Market to Surpass $68 Billion in 2010," *Gartner Research*, June 22, 2010. [Online] Available: http://www.gartner.com/it/page.jsp?id=1389313

[3] D. Krahl, "The Extend simulation environment," *Simulation Conference, 2001. Proceedings of the Winter* , vol.1, pp.217-225, 2001.

[4] N. Leavitt, "Is Cloud Computing Really Ready for Prime Time?," *Computer*, vol. 42, no. 1, pp. 15-20, Jan. 2009.

[5] A. Li, X. Yang, S. Kandula, M. Zhang, "CloudCmp: Shopping for a Cloud Made Easy," in *2nd USENIX Workshop on Hot Topics in Cloud Computing (HotCloud)*, June 2010.

[6] P. Mell, T. Grance "The NIST Definition of Cloud Computing." National Institutes of Standards and Technology [6]. Version 15. 7 October 2009. [Online] Available: http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc

[7] Sun Microsystems. *Introduction to Cloud Computing Architecture*, 1st edition, June 2009.

[8] United States. The Chief Information Officers Council. *Federal Enterprise Architecture Framework*, version 1.1, September 1999. [Online] Available: www.cio.gov/documents/fedarch1.pdf

[9] United States. Federal Chief Information Officer. *State of Public Sector Cloud Computing*. May 20, 2010.

[10] Kaiqi Xiong, Harry Perros, "Service Performance and Analysis in Cloud Computing," services, pp.693-700, *2009 Congress on Services - I*, 2009.

7.0 ACKNOWLEDGMENTS

The authors thank Michael Farber, Michael Cameron, and Paul Brown for funding this study through the Booz Allen Cloud Computing Initiative. The authors also thank Adam Young, Matt Ostrander, and Justin Neroda for their subject matter expertise and guidance.

Booz | Allen | Hamilton

Performance Analysis of Cloud Computing Architectures using Discrete Event Simulation

John Stocker and Andrew Golomb

MODSIM WORLD Conference & Expo

October 2010
Hampton, VA

Motivation

- We know that cloud computing can provide economic benefits by increasing the utilization of server hardware, but we need to know more to effectively apply cloud computing:
  - Can cloud computing also help my organization become more productive, if so by how much?
  - What cloud deployment model should I use?
  - How big should my cloud be if I build a private or community cloud?
  - If I purchase cloud services, what quality of service terms do I need to negotiate into my service level agreement?

To date quantitative analysis of cloud computing has focused on cloud economics and not organizational effectiveness or architecture performance

Cloud Computing Effectiveness Model Framework

A model framework to quantify cloud computing alternatives

- **Evaluate resource use and system performance**
  - Impact of cloud architecture on system performance and reliability
  - Cloud capacity (nodes, processors, memory, storage, etc.)
  - Service response time, capacity, processing and availability
  - Identify network utilization and limiting factors on performance
  - Evaluation of system reliability and effect of resource allocation algorithms
- **What-if analysis**
  - Ability of cloud to meet service level requirements for a range of scenarios
  - Compare private, community, public, and hybrid cloud delivery models
  - Compare the effects of service utilization on response time
  - Trade performance and cost
- **Environment to rapidly design, analyze, and test cloud configurations**
  - Verification of requirements via analysis
  - Provide estimate of expected performance realized through cloud implementation

What is Cloud Computing?

- Many definitions exists, but all definitions share key tenets
  - Massively scalable and elastic IT-enabled infrastructure
  - Delivering capabilities 'as a service' to users or systems
  - Using well-known internet technologies

Cloud Characteristics	Cloud Delivery Models	Cloud Deployment Models
On-demand self-service	Software as a Service (SaaS)	Private Cloud
Ubiquitous network access	Platform as a Service (PaaS)	Community Cloud
Massive scale	Infrastructure as a Service (IaaS)	Public Cloud
Rapid elasticity		Hybrid Cloud
Pay per use		

Source: NIST working definition

Cloud Deployment Model—Private Clouds

Private Clouds	Public Clouds
Community Clouds	Hybrid Clouds

- Only leverages internal cloud infrastructure
- Organization buys and maintains all cloud infrastructure
- Often used without multi-tenancy or virtualization

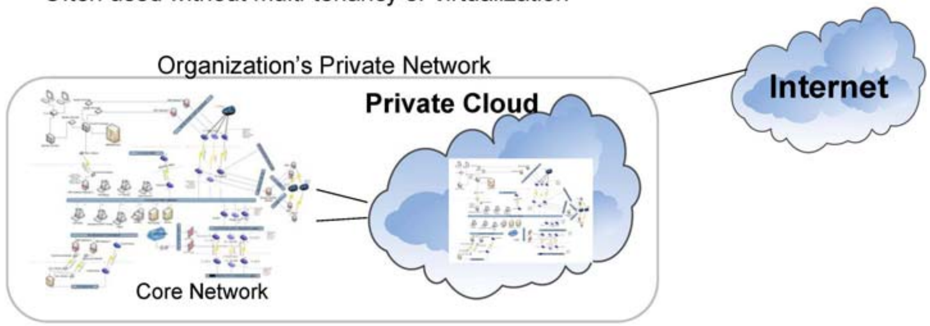

MODSIM WORLD Conference & Expo

Booz | Allen | Hamilton 4

Cloud Deployment Model—Community Clouds

Private Clouds	Public Clouds
Community Clouds	Hybrid Clouds

- Cloud infrastructure is shared by several organizations
- The cloud is tailored for a specific set of missions
- Executive agent can be assigned for governance

MODSIM WORLD Conference & Expo

Booz | Allen | Hamilton 5

Cloud Deployment Model—Public Clouds

Private Clouds	**Public Clouds**
Community Clouds	Hybrid Clouds

- Cloud infrastructure is owned by an organization selling services
- Organizations lease time & space in the cloud
- Highly virtualized with high multi-tenancy
- This is the predominant model for commercial cloud computing

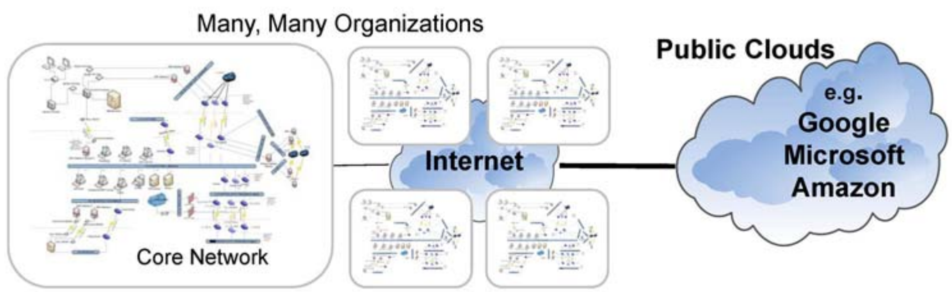

Cloud Deployment Model—Hybrid Clouds

Private Clouds	Public Clouds
Community Clouds	**Hybrid Clouds**

- Allows for Cloud Bursting – only using public clouds as needed
- Organization buys and maintains most of the cloud infrastructure

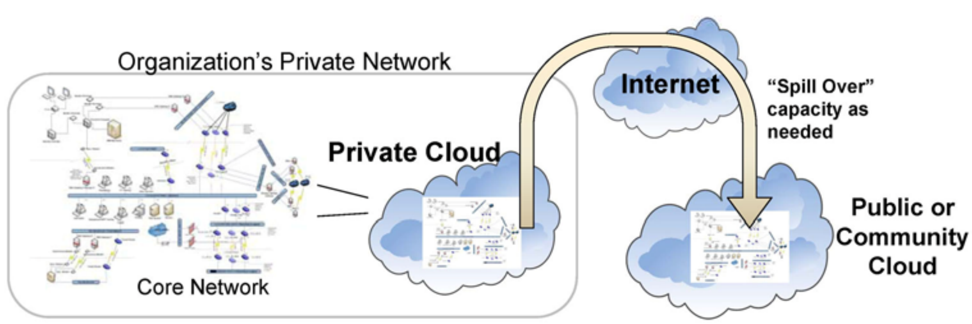

Cloud Computing is not a Panacea

- **Cloud Computing Readiness Assessment & Transition Models**–not everything belongs in the cloud
  - Real-time applications, command & control systems are probably not good candidates for cloud environments
- **Economics**
  - Trade space among capital expenditures, operational expenditures, migration, and hidden costs
- **Legal & Regulatory Challenges**
  - HIPAA, SOX, privacy, physical data may be located outside US
- **Multi-Tenet Environment**
  - All resources are shared (disk storage, network) causes difficulty cleaning ‘spills’
- **Security**
  - Can security actually be *increased* by moving to a cloud?
  - What new security vulnerabilities are introduced?
- **Technical Maturity**
  - Cloud computing is proven at Google, Microsoft, Amazon, using proprietary technology–open-source cloud technologies are maturing

Why Discrete Event Simulation?

- Discrete event simulation combines business process reengineering with proven modeling techniques to determine the potential “to be” process performance, efficiency, and variation
  - Process models—abstract representations of processes—can be developed and analyzed prior to investing in process changes
  - Advances in simulation and computer technology allow for rapid development and execution of scenarios using models
  - A limitless amount of analysis can be executed to investigate the affect of process or architecture changes on performance
  - Approach provides quantitative metrics on performance to serve as justification of investment in resources and process modifications

Discrete Event Simulation Characteristics
Dynamic: evolves over time
Stochastic: randomness introduced for variables
Discrete time: significant process changes occur at instances in time

Cloud Computing Performance Model

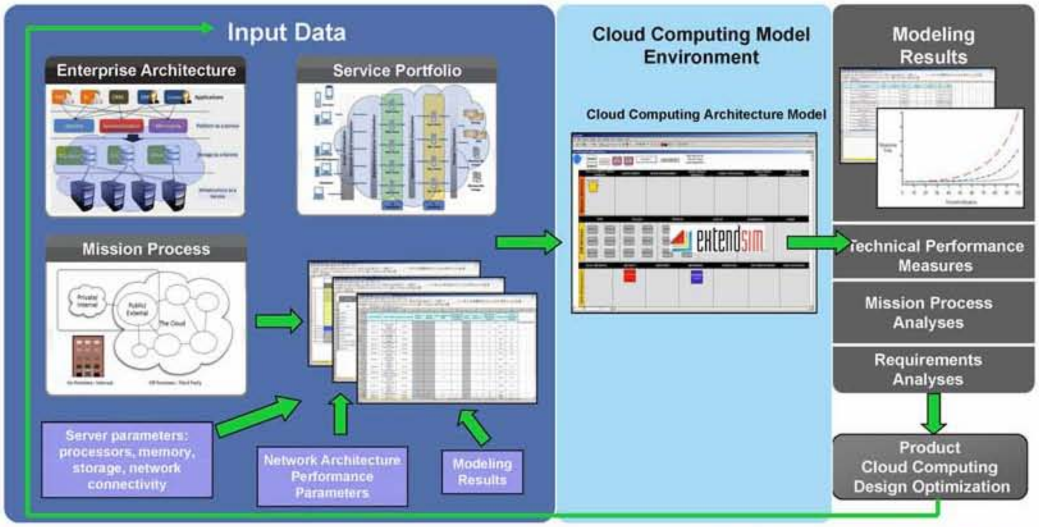

Business Process Model

Office site name
Number of staff members
Communication link characteristics - Maximum uplink bandwidth (Mbps) - Maximum downlink bandwidth (Mbps)

- Workforce model
  - The number of staff members per office site remains constant
  - Differences in application and service usage between staff members based on their roles is aggregated in the application request rate
  - Uplink and downlink bandwidth between the office and the cloud
- Application and service model
  - Request rates vary over time: business day, weekly, monthly, quarterly, etc.
  - Computational requirements and data transferred are characterized by statistical distributions with application specific parameters

Application or service name
Request rate (requests over time)
Data transferred (KB)
Computational requirements (MIPS)
QoS requirements - Allowable response latency (seconds)

Cloud Computing Architecture Model

- Cloud deployment model
  - Clouds are collections of compute nodes—individual servers networked together
  - Nodes within the cloud are assumed to be homogenous—using homogeneous nodes is an industry best practice to reduce maintenance costs
  - Generated service requests are assigned to the least utilized nodes
  - If the maximum process capacity is exceeded, requests are then delayed through multi-tasking

Cloud name
Number of Nodes
Node processing capacity (MIPS / sec)
Node memory (GB)
Node storage capacity (TB)
Node storage access speed (milliseconds)
Inter-node network characteristics

Sample Analysis—Input Data

Office Characteristics

- Single office

Parameter	Value
Number of staff members	5,000
Maximum Uplink Bandwidth (Mbps)	?
Maximum Downlink Bandwidth	?

Cloud Architecture

- Private Cloud

Parameter	Value
External Network Bandwidth (Mbps)	?
Number of nodes	?
CPU Processor speed (GHz)	2
Instructions per cycle	16
User storage capacity per node (TB)	0.1
RAM (GB)	1

Sample Analysis—Application and Service Characteristics

	Daily usage per user		Data Sent/Received per Request (KB)		Allowable Latency (seconds)	Million CPU Instructions Required per Request	
Service Request	Average	Maximum	Average	Maximum		Average	Maximum
Email	30	500	4	1000	2	0.1	10
Database access	3	5	3000	60000	120	1000	10000
Instant messaging	50	500	1	2	2	0.05	1
Email w/ attachment	5	20	2000	20000	60	1	10
Wiki/Blog entry	1	4	2	40	5	1	10
Office app usage	30	100	8	500	1	10	100
Model / simulation	8	16	200	1000000	2400	100000	10000000
Financial Report Gen	1	3	10	40	1200	10000	100000
Video chat	5	20	6	40	1	1000	10000
Collaboration site	10	20	20	100000	180	10	100

– Triangular statistical distribution models are used

Sample Bandwidth Sensitivity Analysis

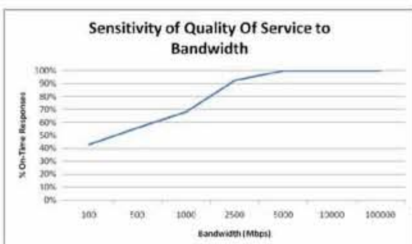

Assumptions:

- Users: 5,000
- Number of Nodes: 500
- Uplink and downlink rates are identical (synchronous link)

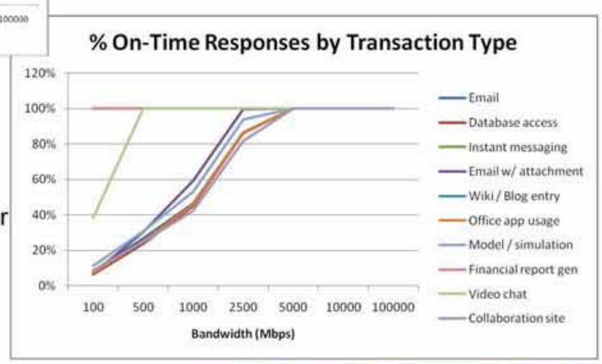

Conclusion:

~ 1 Mbps of bandwidth / user is required to provide high quality of service, especially for highly interactive applications (e.g. video chat)

Sample Node Count Sensitivity Analysis

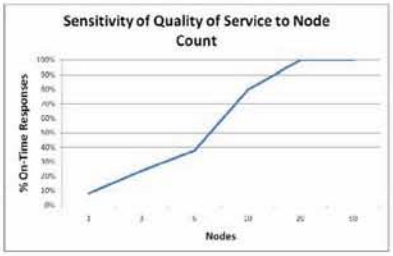

Assumptions:
- Users: 5,000
- Bandwidth: 5 Gbps
- Uplink and downlink rates are identical (synchronous link)
- 0.1 TB user storage per node

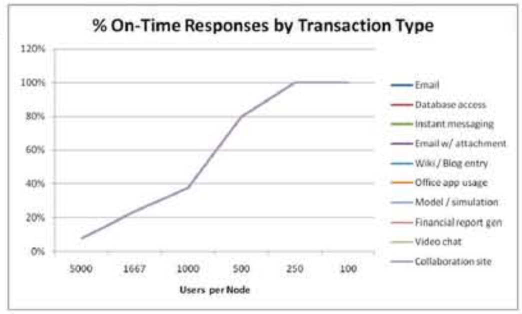

Conclusion:

~ 20 nodes, or 250 users per node is required to provide high quality of service

MODSIM WORLD Conference & Expo

Sample Node Count Sensitivity Observations

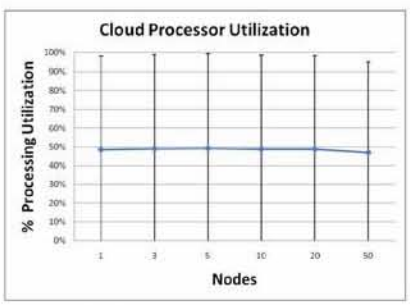

- Peak processor utilization slowly declines as more nodes are added
- Peak storage utilization quickly declines with the addition of more nodes

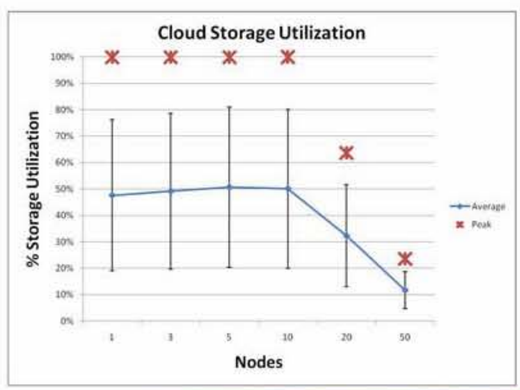

Conclusion:
- Performance in sample is limited based on storage
- Should consider increasing user storage per node instead of adding more nodes based on cost

MODSIM WORLD Conference & Expo

Conclusion and Future Work

- We have developed a model framework for analyzing the performance of candidate cloud computing architectures and resultant organizational effectiveness
- Baseline future models using data from Booz Allen's cloud computing infrastructure that supports 30,000+ employees
- Business process modeling can be improved through integrating industry standard process model data captured in FEAF or DoDAF artifacts
- Coupling quantitative performance and a cost models will provide an overarching analytical approach to develop cloud migration strategies
- Detailed modeling and analysis of cloud computing architectures is necessary to capture trades among:
  - Cloud deployment models: private, community, public, hybrid
  - Fewer more powerful nodes vs. many less powerful nodes
  - Commodity hardware vs. more expensive, high-availability hardware

Questions and Comments

John Stocker
Lead Associate
Booz | Allen | Hamilton
Booz Allen & Hamilton Inc.
8283 Greensboro Drive
Mclean, VA 22102
Tel (703) 377 6773
Mobile (571) 403 0613
Fax (703) 902 3392
stocker_john@bah.com

Andrew Golomb
Senior Consultant
Booz | Allen | Hamilton
Booz Allen & Hamilton Inc.
8283 Greensboro Drive
Mclean, VA 22102
Tel (703) 377 8641
Fax (703) 902 3392
golomb_andrew@bah.com