

Cloud Computing for Business

The Open Group Guide



Chris Harding and members of
The Open Group Cloud Work Group

Cloud Computing for Business - The Open Group Guide

Other publications by Van Haren Publishing

Van Haren Publishing (VHP) specializes in titles on Best Practices, methods and standards within four domains:

- IT and IT Management
- Architecture (Enterprise and IT)
- Business Management and
- Project Management

Van Haren Publishing offers a wide collection of whitepapers, templates, free e-books, trainer materials etc. in the **Van Haren Publishing Knowledge Base**: www.vanharen.net for more details.

Van Haren Publishing is also publishing on behalf of leading organizations and companies: ASLBiSL Foundation, BRMI, CA, Centre Henri Tudor, Gaming Works, IACCM, IAOP, Innovation Value Institute, IPMA-NL, ITSqc, NAF, Ngi/NGN, PMI-NL, PON, The Open Group, The SOX Institute.

Topics are (per domain):

IT and IT Management

ABC of ICT
ASL®
CATS CM®
CMMI®
COBIT®
e-CF
ISO 20000
ISO 27001/27002
ISPL
IT-CMF™
IT Service CMM
ITIL®
MOF
MSF
SABSA

Architecture (Enterprise and IT)

ArchiMate®
GEA®
Novius Architectuur Methode
TOGAF®

Business Management

BABOK® Guide
BiSL®
BRMBOK™
EFQM
eSCM
IACCM
ISA-95
ISO 9000/9001
Novius B&IP
OPBOK
SAP
SixSigma
SOX
SqEME®

Project, Program and Risk Management

A4-Projectmanagement
DSDM/Atern
ICB / NCB
ISO 21500
MINCE®
M_o_R®
MSP™
P3O®
PMBOK® Guide
PRINCE2®

For the latest information on VHP publications, visit our website: www.vanharen.net.

Cloud Computing for Business

The Open Group Guide

THE
Open
GROUP



Colofon

Title:	Cloud Computing for Business - The Open Group Guide
A Publication of:	The Open Group
Lead Author:	Chris Harding
Editorial Committee:	Pamela K. Isom, Mark Skilton, and Chris Harding
Publisher:	Van Haren Publishing, Zaltbommel, www.vanharen.net
ISBN:	978 90 8753 657 2
Edition:	First edition, first impression, August 2011 First edition, second impression, November 2015
Design and Layout:	CO2 Premedia bv, Amersfoort – NL
Copyright:	© The Open Group, 2011

For any further enquiries about Van Haren Publishing, please send an e-mail to:
info@vanharen.net

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of the copyright owner.

The views expressed in this document are not necessarily those of any particular member of The Open Group.

It is fair use of this specification for implementers to use the names, labels, etc. contained within the specification. The intent of publication of the specification is to encourage implementations of the specification.

This specification has not been verified for avoidance of possible third-party proprietary rights. In implementing this specification, usual procedures to ensure the respect of possible third-party intellectual property rights should be followed.

Comments relating to the material contained in this document may be submitted to:

The Open Group
Apex Plaza, Forbury Road
Reading
Berkshire RG1 1AX
United Kingdom
or by electronic mail to:
ogspecs@opengroup.org

Contents

Preface.....	XIII
About The Open Group	XXII
Trademarks.....	XXIII
Acknowledgements	XXIV
Referenced documents	XXVIII
1 What is cloud?	1
1.1 NIST definition of cloud computing.....	2
1.2 Essential characteristics	4
1.2.1 On-demand self-service	4
1.2.2 Broad network access.....	5
1.2.3 Resource pooling.....	6
1.2.4 Rapid elasticity.....	9
1.2.5 Measured service	10
1.3 Service models.....	11
1.3.1 Cloud infrastructure as a service (IaaS)	12
1.3.2 Cloud platform as a service (PaaS)	13
1.3.3 Cloud software as a service (SaaS)	14
1.4 Deployment models.....	15
1.4.1 Private cloud.....	15
1.4.2 Public cloud.....	16
1.4.3 Community cloud	16
1.4.4 Hybrid cloud	17
1.5 Extending the NIST model to business processes.....	18
1.6 Comparison with traditional environments	18
1.7 Providing and using cloud services	20
1.7.1 Providing cloud services.....	20
1.7.2 Providing added services.....	22
1.7.3 Developing added services.....	23
1.7.4 Using cloud services.....	25
1.8 The impact of cloud computing.....	27
1.8.1 New business paradigm.....	28
1.8.2 Cloud ecosystems	30

2	Why cloud?	33
2.1	Agility.....	35
2.1.1	Changing business processes	35
2.1.2	Development and testing.....	36
2.1.3	Resource scaling	36
2.1.4	Reduced need for training.....	37
2.2	Productivity.....	38
2.2.1	Collaborative working	38
2.2.2	Shared logic	38
2.3	Quality	39
2.3.1	Better usage information	39
2.3.2	Better manageability	40
2.3.3	Better quality of IT provision.....	40
2.3.4	Better business continuity	41
2.3.5	Better carbon footprint	41
2.4	Cost.....	42
2.4.1	Server consolidation.....	42
2.4.2	Thin clients	44
2.4.3	Community cost sharing.....	44
2.4.4	Replacing CAPEX with OPEX.....	44
2.5	New business opportunities	45
2.5.1	Cloud service provision	45
2.5.2	Added service provision	46
3	Establishing your cloud vision	47
3.1	Understanding the business context	48
3.1.1	Basic situation	48
3.1.2	Business goals.....	48
3.1.3	Risk-reward balance.....	48
3.1.4	Impact on products and services.....	50
3.1.5	Business processes.....	51
3.1.6	Scope and complexity	53
3.1.7	Collaboration <i>versus</i> information restriction.....	54
3.2	Three example cloud projects	55
3.2.1	Konsort-Prinz	55
3.2.2	Sam Pan Engineering.....	56
3.2.3	ViWi	57

3.3	Assessing cloud suitability – the cloud buyer’s decision tree	57
3.3.1	Assessment considerations.....	58
3.3.2	Question 1: Is your business situation vertical?	59
3.3.3	Question 2: Are the processes differentiating?	60
3.3.4	Question 3: Are there impediments to outsourcing?.....	60
3.3.5	Question 4: Are there impediments to cloud adoption?	61
3.3.6	Question 5: Is the primary business driver cloud-compatible?	62
3.3.7	Question 6: Will the solution be a platform?	64
3.3.8	Question 7: Is the application insulated from changes to the business process?	66
3.3.9	Question 8: Is the differentiation IT-based?.....	67
3.3.10	Question 9: Are the hardware, operating system, and application custom-made?	67
3.3.11	Question 10: Are the hardware and operating system custom- made or specialized?	68
3.4	The example project visions.....	69
3.4.1	Konsort-Prinz	70
3.4.2	Sam Pan Engineering.....	72
3.4.3	ViWi	73
4	Buying cloud services	75
4.1	Determining fit	76
4.1.1	Workload and cost models.....	77
4.1.2	Workload factors	77
4.1.3	Workload allocations	78
4.1.4	Pay-as-you-go <i>versus</i> ownership	80
4.1.5	Modeling resources and costs.....	81
4.1.6	Example workload and cost models – Konsort-Prinz.....	82
4.1.7	Example workload and cost models – ViWi.....	83
4.1.8	Using the models	84
4.2	Establishing requirements	85
4.2.1	Service functionality	86
4.2.2	Back-up.....	87
4.2.3	Bulk data transfer	88
4.2.4	Supplier choice.....	88
4.2.5	Availability.....	89

4.2.6	Reliability.....	90
4.2.7	Recoverability.....	90
4.2.8	Responsiveness	91
4.2.9	Throughput.....	93
4.2.10	Configurability.....	93
4.2.11	Reporting.....	94
4.2.12	Fault management	94
4.2.13	End user access control.....	94
4.2.14	Provider access control	95
4.2.15	Resource partitioning	96
4.2.16	Logging	96
4.2.17	Threat management.....	96
4.2.18	Compliance with regulations	97
4.3	Selection.....	98
4.3.1	Exit strategy.....	98
4.3.2	Contract terms.....	100
4.3.3	Negotiation.....	101
4.3.4	Choosing the service.....	101
4.4	Monitoring	103
4.4.1	Workload and cost.....	103
4.4.2	Conformance to requirements.....	105
5	Understanding cloud risk	109
5.1	Risk management	110
5.1.1	Risk assessment.....	110
5.1.2	Risk communication	111
5.2	Cloud mission risks	113
5.2.1	Financial	115
5.2.2	Organization and culture	115
5.2.3	Service integration.....	116
5.2.4	Compliance	118
5.2.5	Business continuity management.....	118
5.2.6	System quality.....	119
5.2.7	External service.....	120
5.3	System quality risk factors.....	122
5.3.1	Functionality	122
5.3.2	Performance	122
5.3.3	Manageability.....	124

5.3.4	Security	124
5.3.5	User satisfaction.....	125
5.4	Continuing risk assessment	125
5.4.1	Solution architecture development	126
5.4.2	Cloud service selection and procurement.....	128
5.4.3	Solution operation.....	129
6	Building ROI from cloud computing	131
6.1	Productivity: more business with less IT	132
6.1.1	Resource utilization.....	133
6.1.2	Usage-based pricing	133
6.1.3	Specialization and scale	135
6.2	Speed: getting there more quickly.....	135
6.2.1	Time to deployment	135
6.2.2	Lifetime cost models	136
6.2.3	IT asset management	137
6.3	Size: breaking new ground.....	138
6.3.1	Entering new markets	138
6.3.2	High-value services	138
6.3.3	The long tail.....	139
6.3.4	Becoming a cloud provider	139
6.4	Quality: improved margin from better service.....	140
6.4.1	Competitive pressure	140
6.4.2	The importance of quality	141
6.5	Comparing ROI of cloud and traditional IT solutions.....	141
6.5.1	Basis of ROI calculations	142
6.5.2	Konsort-Prinz	143
6.5.3	Sam Pan Engineering	144
6.5.4	ViWi	145
6.6	Measuring and tracking ROI.....	147
6.6.1	Utilization	147
6.6.2	Time compression	148
6.6.3	Scale.....	148
6.6.4	Quality	148
6.6.5	Examples.....	149
6.6.6	Konsort-Prinz	149
6.6.7	Sam Pan Engineering.....	151
6.6.8	ViWi	153

7	The challenge	155
Appendix A	Cloud computing in use	159
A1	Cross-industry use-cases.....	160
A2	Financial services use-cases.....	164
A3	Government use-cases.....	167
A4	Telecommunications operator use-cases	171
A5	Media and entertainment use-cases	172
A6	Health services use-cases	173
A7	Pharmaceuticals use-cases	174
A8	Distribution use-cases	175
A9	Energy and utilities use-cases.....	176
A10	Higher education use-cases	177
A11	Use-case actors	178
A12	Use-case benefits	184
Appendix B	Glossary.....	195
Index.....		211

Table of figures

Figure 1.1: Basic cloud computing model.....	3
Figure 1.2: The essential characteristics of cloud computing	4
Figure 1.3: Cloud service models	12
Figure 1.4: Providing cloud services	21
Figure 1.5: Providing added services	22
Figure 1.6: Using cloud services	26
Figure 1.7: Use of cloud services to support business processes	29
Figure 1.8: Internal and external cloud services.....	29
Figure 1.9: Business ecosystem	30
Figure 1.10: Cloud ecosystem	31
Figure 2.1: Use-case business elements	33
Figure 3.1: Risk <i>versus</i> reward.....	49
Figure 3.2: Impact on products and services	50
Figure 3.3: Business operating scope <i>versus</i> complexity	53
Figure 3.4: Collaboration <i>versus</i> information restriction.....	54
Figure 3.5: Deciding on your cloud solution	58
Figure 3.6: Consort-Prinz decision tree.....	70
Figure 3.7: Sam Pan Engineering decision tree	72
Figure 4.1: The cloud buying lifecycle.....	75
Figure 4.2: Types of workload allocation.....	79
Figure 4.3: Examples of workload allocation types.....	79
Figure 4.4: Example fixed workload costs	81
Figure 4.5: Example variable workload costs.....	81
Figure 4.6: ViWi workload model	84
Figure 4.7: ViWi cost models	103
Figure 4.8: Consort-Prinz capacity utilization.....	104
Figure 4.9: Consort-Prinz peak day utilization	104
Figure 4.10: Consort-Prinz first year costs.....	105
Figure 4.11: Performance and manageability comparison.....	106
Figure 4.12: Comparison of actual performance with SLA	107
Figure 5.1: ViWi risk exposure	113
Figure 5.2: Sam Pan Engineering initial risk assessment	114
Figure 5.3: ViWi initial system quality assessment	120
Figure 5.4: ViWi external cloud service system quality assessment	121
Figure 5.5: Consort-Prinz initial risk assessment.....	126
Figure 5.6: Consort-Prinz risk mitigation.....	127

Figure 5.7: Konsort-Prinz supplier cost models.....	128
Figure 5.8: ViWi six-month risk assessment.....	129
Figure 6.1: The capacity <i>versus</i> utilization curve.....	132
Figure 6.2: Software license costs	134
Figure 6.3: Software cost optimization	134
Figure 6.4: Time to deployment	136
Figure 6.5: Speed of cost reduction	137
Figure 6.6: Revenue, cost, and margin.....	140
Figure 6.7: Competitive pressure.....	141
Figure 6.8: Konsort-Prinz data center upgrade ROI.....	143
Figure 6.9: Konsort-Prinz cloud solution ROI	144
Figure 6.10: Sam Pan Engineering ROI calculation.....	145
Figure 6.11: ViWi Years 1 and 2 cumulative revenue and costs	146
Figure 6.12: ROI for the backers of ViWi.....	146
Figure 6.13: Konsort-Prinz November ROI scorecard.....	150
Figure 6.14: Sam Pan Engineering six-month ROI scorecard.....	151
Figure 6.15: Sam Pan Engineering reduced scale ROI projection	152
Figure 6.17: ViWi ROI projection after six months.....	153
Figure 6.16: ViWi six-month ROI scorecard	153

Preface

This Guide is about how to use cloud computing to gain business benefit for your enterprise.

Cloud computing became a hugely popular topic in 2009. Wherever you looked on the web, or in IT organizations' marketing and advertising material, there were references to 'the cloud', or services in or from 'the cloud'. Everything was 'as a service'. The initial excitement has now subsided, but it is clear that cloud computing is a very important technological development. Companies have started to use it in earnest, and will do so increasingly for the foreseeable future. It is time to take an in-depth look at what cloud computing is and, more particularly, how to take advantage of its potential.

From your own personal experience you may think you know how good cloud computing is. Perhaps you use a web mail service or one of the popular social networking sites, or you download applications to your phone from an application store for very little cost and as easy as a couple of clicks. Perhaps you back your PC files up to a service somewhere in 'the cloud'.

There are many reasons why you personally can benefit from cloud computing – but what about your company? Many enterprises are now thinking about moving IT services into the cloud. What's in it for them?

Here are some of the reasons why companies are using, or thinking of using, cloud computing that were given in business scenario workshop sessions held at The Open Group conferences:

“We are unable to align capabilities with the needs of the business.”

“We need rapid access to different, and potentially game-changing, models of computing and new technologies. Without this, we could be left in the dust by competition.”

“I don't want to invest in capital in very early stages – I want to wait until there are signs that the business will survive.”

“There are huge risks associated with storing data that we aren’t competent to manage. We want to get rid of it but still access it through a competent authority.”

“Our current environment is overly-dependent on a few key individuals to support technology – a huge risk.”

“We have no physical space – we’re full!”

There are many other reasons too. The benefits of moving to the cloud are categorized later in this Guide as agility, productivity, quality, cost, and new business opportunities. Can cloud computing meet all of these expectations? Perhaps it can, but there is still much confusion about what it is, what it can do, and how to use it.

Why this Guide is important

Every few years there is a new hot topic that becomes the focus for media and analyst attention. Enormous benefits are stated, and there are staggering predictions for market growth. The benefits may be real, and the predictions may be accurate, but they are given in such general terms that it is hard for a particular company to understand how it can take advantage of the new phenomenon and share in the growth.

This is certainly the case with cloud computing. Many benefits are claimed, including greater agility, lower cost, improved security, reduced risk, easier compliance with regulation, higher-quality IT support, and better business continuity. Analysts predict compound annual growth rates of over 25% in the cloud computing market.

“So does this mean,” you might ask, “that my company can grow by 25% by using cloud computing? And, if so, how exactly do we do it?”

The authors of this Guide believe that cloud computing is a major development in IT, that it will grow as predicted, and that it can deliver real business benefits to companies.

It is a complex concept. It is not based on a single technological breakthrough, but comes about through the combination of several innovations and improvements, most notably the development of virtualization, the increasing capacity of the Internet and the growing sophistication of Internet-based technologies. It has five essential characteristics: on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service. It has three service models: Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). It has four deployment models: private cloud, public cloud, community cloud, and hybrid cloud.

Essential Characteristics	Service Models	Deployment Models
On-demand self-service Broad network access Resource pooling Rapid elasticity Measured service	Infrastructure as a Service (IaaS) Platform as a Service (PaaS) Software as a Service (SaaS)	Private cloud Public cloud Community cloud Hybrid cloud

Table 1: The characteristics and models of cloud computing

You cannot simply say: “We will use cloud computing,” and expect your company to be transformed, as at the wave of a magic wand. It may not be the best choice in every situation, or in every enterprise. You must understand cloud computing, the different forms it can take, the different ways it can be used, and the different ways in which it can benefit your company.

For a company to adopt cloud computing successfully, depending on the change it is imposing, various stakeholders must be involved. As is also the case with traditional outsourcing, new skills may be needed and old skills may be obsolete. The transition may imply a large change-management exercise.

The company’s executives must have a shared understanding. Many parts of the company are affected, and there are complex decisions to be made. Making the transition to cloud computing requires a corporate team effort.

The Open Group is a consortium of companies that provide and use IT products and services. Many of those companies are thinking of using cloud computing, some are using cloud computing, and some are providing cloud

computing solutions. They have gained or are gaining the understanding that they need for this. This Guide does not represent the views of any one of those companies, but it reflects their combined experience, and makes their shared understanding available to the wider IT community.

Cloud computing has a huge potential. To realize that potential, a company must use the form of cloud computing most suited to its needs, and in the way that will give it the best advantage. The aim of this Guide is to give you the information and understanding that you need to help your company to do this, and gain the maximum benefit from cloud computing.

Who this Guide is for

Cloud computing will affect many people in any organization that uses it, both directly and indirectly.

Members of the governing board of an enterprise may not be concerned with day-to-day operations, but will need to understand the language of cloud computing and the key factors involved in its use in order to provide overall direction.

Cloud computing has significant impact on IT investment and running costs. It also changes the risks that the organization runs in relation to IT. The CIO or CTO is responsible for ensuring that an organization has the IT capabilities that it needs, and that they give a good return on the money invested in them. The ROI is also of prime concern to the CFO. All of these officers are concerned with the risks run by the organization. The use of cloud computing is likely to have an impact on the enterprise's ability to change and innovate, risk model, financial model, and organizational structure. This impact concerns them, and should also concern the CEO.

IT managers are responsible for the provision and operation of the IT capabilities, and for the management of the departments that support them. They will be intimately involved in the use of cloud computing.

Line-of-business managers are responsible for the effective operation of their departments, and for making a profit or using resources cost-effectively. They will be concerned to ensure that cloud computing increases effectiveness and

profitability in their business areas. In non-commercial enterprises, heads of departments responsible for service delivery will have similar concerns.

Marketing managers will be concerned with the reduced time-to-market that cloud computing can bring, and in the impact that its use might have on the enterprise's image.

Many large enterprises have specialist risk managers. They will be concerned with the cloud computing risk model. Security managers will be concerned with the different threats that are presented by the adoption of cloud computing. Compliance managers will be concerned with whether the use of cloud computing is affected by laws or other regulations to which the enterprise is subject.

Procurement managers will be concerned with choosing and buying cloud services, and with gaining feedback on the value of those services in operation.

Business and enterprise architects want to understand cloud from a business point of view. IT consultancies providing outsourcing or integrating cloud services on behalf of organizations need to know how they can best deliver the services their clients require.

This Guide is for all of these people, and indeed for all executives whose companies are using, or thinking of using, cloud computing.

How this Guide is organized

This Guide is organized as described below.

Chapter 1: What is cloud?

The first chapter explains what cloud computing is and describes how it can be used.

The explanation is based on the standard definition of cloud computing developed by the US National Institute of Science and Technology [NIST]. It includes an overview of the definition (Section 1.1) and sections on cloud

computing Essential Characteristics (Section 1.2), Service Models (Section 1.3), and Deployment Models (Section 1.4).

This is followed by a brief discussion of an extension of the NIST model that many people advocate, Business Process as a Service (Section 1.5), and a summary comparison with traditional environments (Section 1.6).

Organizations can take different roles in providing cloud services, developing and providing added services that are based on cloud services, and using cloud services and added services. There are different issues and considerations that impact each role. This is discussed in Providing and Using Cloud Services (Section 1.7).

The really exciting thing about cloud computing is its potential to transform business. The final section of the chapter (Section 1.8) sketches this broader impact. Organizations can create and use IT and business services on-demand from optimal sources to maximize utilization and cost-effectiveness, and this is leading to a new paradigm with cloud-supported business ecosystems.

Chapter 2: Why cloud?

This chapter discusses the main reasons why enterprises are adopting cloud computing, under the headings of Agility (Section 2.1), Productivity (Section 2.2), Quality (Section 2.3), Cost (Section 2.4), and New Business Opportunities (Section 2.5).

Chapter 3: Establishing your cloud vision

Your business situation is either a problem or an opportunity for which you are seeking a solution that includes IT enablement. You see a technological possibility as the way to solve your problem, or seize your opportunity. This is your architecture vision. This chapter describes how to establish an architecture vision with cloud computing as the technical possibility.

The first essential step in establishing the vision is to ensure that you understand your business context. The chapter starts by putting forward a set of considerations that will help you achieve that understanding (Section 3.1). This section is derived from the Cloud Buyers Requirements Questionnaire [BUYERSQ] previously published by The Open Group.

The book includes three large examples to illustrate and explain its ideas. These describe fictional companies that are chosen to show different aspects of cloud computing. These examples are introduced in the next section of the chapter (Section 3.2). Spreadsheets containing the figures and calculations from these examples are available for download from The Open Group web site [EXAMPLES].

There follows a cloud buyer's decision tree (Section 3.3) that helps you to determine whether cloud computing belongs in your vision, and in what form. Cloud computing can bring business benefits to many enterprises. It is, however, not always the right solution, and there are several forms of cloud computing, which have different advantages in different situations. The decision tree is an aid to decision-making, not a definitive algorithm, and was previously published by The Open Group as a White Paper [BUYERSTREE].

The final section of the chapter (Section 3.4) illustrates the application of the decision tree to establish a cloud vision, using the three large examples.

Chapter 4: Buying cloud services

Many enterprises that adopt cloud computing will do so by buying cloud services, and the process of doing so is described in this chapter. The chapter includes discussion of the key factors to consider – cost models, security, availability, performance, manageability, and so on – and describes how to model the use of cloud services. This material is assumed by some parts of the subsequent chapters, and is important even if you will not be directly involved in purchasing a cloud service.

Buying and using cloud services follows a lifecycle, in which services are selected and their use is reviewed as a basis for renewing or replacing them. The chapter follows this lifecycle, describing the successive phases of Determining Fit (Section 4.1), Establishing Requirements (Section 4.2), Selection (Section 4.3), and Monitoring (Section 4.4).

Chapter 5: Understanding cloud risk

Risk is a fundamental issue for any enterprise. This chapter is about how to understand the main risks associated with cloud computing. Its approach is based on the Mosaic approach to risk management that was developed at the Carnegie-Mellon University Software Engineering Institute (CMU SEI).

Risk management is a core business activity of all enterprises, large and small. The first section (Section 5.1) gives a brief introduction to the Mosaic approach to the topic.

In this approach, a *mission risk* is a systemic risk that affects a program's ability to achieve its key objectives. The chapter describes the typical mission risks for cloud computing projects (Section 5.2), and explains the particularly complex factors that affect one of them – system quality (Section 5.3).

The final section of the chapter describes the process of risk assessment for cloud computing projects (Section 5.4).

Chapter 6: Building ROI from cloud computing

Return on Investment (ROI) is perhaps the most widely-used measure of financial success in business. If you have a proposal to use cloud computing in place of in-house IT, this is how you and others will want to assess it. This chapter discusses the qualities of cloud computing that affect ROI, and describes how to measure, and maximize, your ROI from cloud computing. It is based on the White Paper Building Return on Investment from Cloud Computing [CLOUDROI] previously published by The Open Group.

How does cloud computing contribute to ROI? There are a number of fundamental drivers that impact on investment, revenue, cost, and timing that can be positively influenced by using cloud services. They are described in the first four sections of the chapter. They relate to productivity (Section 6.1), speed (Section 6.2), size (Section 6.3), and quality (Section 6.4).

The remaining two sections describe how to compare cloud and traditional IT solutions (Section 6.5), and how to monitor the drivers to maintain and build ROI from cloud computing (Section 6.6). They are illustrated using the three large examples.

Chapter 7: The challenge

The final chapter summarizes the key drivers and issues for cloud computing, and describes the challenge that it presents to the business executive.

Appendix A

The appendix contains the set of cloud business use-cases developed by The Open Group. The set was originally published as a White Paper [CBUC]. The appendix includes some additional material: analyses of the actors in the use-cases, and of the benefits that the use-cases show.

Appendix B

There is also a Glossary of abbreviations and terms used, followed by a list of References.

About The Open Group

The Open Group is a vendor-neutral and technology-neutral consortium, whose vision of Boundaryless Information Flow™ will enable access to integrated information within and between enterprises based on open standards and global interoperability. The Open Group works with customers, suppliers, consortia, and other standards bodies. Its role is to capture, understand, and address current and emerging requirements, establish policies, and share best practices; to facilitate interoperability, develop consensus, and evolve and integrate specifications and Open Source technologies; to offer a comprehensive set of services to enhance the operational efficiency of consortia; and to operate the industry's premier certification service.

Further information on The Open Group is available at www.opengroup.org.

The Open Group has over 15 years' experience in developing and operating certification programs and has extensive experience developing and facilitating industry adoption of test suites used to validate conformance to an open standard or specification.

The Open Group publishes a wide range of technical documentation, the main part of which is focused on development of Technical and Product Standards and Guides, but which also includes White Papers, Technical Studies, and Business Titles.

A catalog is available at www.opengroup.org/bookstore.

Trademarks

Boundaryless Information Flow™ is a trademark and ArchiMate®, Jericho Forum®, Making Standards Work®, Motif®, OSF/1®, The Open Group®, TOGAF®, UNIX®, and the “X” device are registered trademarks of The Open Group in the United States and other countries.

The Open Group acknowledges that there may be other brand, company, and product names used in this document that may be covered by trademark protection and advises the reader to verify them independently.

Acknowledgements

Much of the material in this Guide was derived from the work of the Cloud Computing Explained, Cloud Business Use-Cases, and Cloud Business Artifacts project of The Open Group Cloud Computing Work Group. The Cloud Computing Explained project was led by Gill Hughes of Capgemini and Shuvanker Ghosh of IBM, and included a significant contribution from Ed Harrington of Architecting-the-Enterprise. The Cloud Business Use-Cases project was led by Pamela K. Isom of IBM and Deborah Hawley of The MITRE Corporation, and included notable contributions from Sreekanth Iyer of IBM. The Cloud Business Artifacts project was led by Penelope Gordon, now of 1Plug Corporation, and Mark Skilton of Capgemini.

The material was assembled and rewritten for the Guide under the supervision of an editorial committee consisting of Pamela K. Isom, Mark Skilton, and Chris Harding of The Open Group, who acted as lead author.

The results were reviewed twice, first by a panel of reviewers with particular interest in or experience of the application of Cloud Computing to business situations, including people external to The Open Group, and then by the member companies of The Open Group, as part of The Open Group procedure for approval of publication. These reviews contributed significantly to the quality of the end result.

The following people provided input in the first review. Their independent expertise was extremely valuable, and their comments resulted in a significant restructuring of the early chapters.

- Stuart Boardman, Getronics Consulting
- Rick Crosby, UK Cabinet Office
- Leo Geubbels, Netherlands Ministry of Economic Affairs, Agriculture and Innovation
- Erik Hoekx, Royal Bank of Scotland
- Eddie Michiels, Ernst & Young
- Kevin Holland, UK National Health Service
- Paul Leunissen, Netherlands Ministry of Infrastructure and the Environment

- Karin Middeljans, Netherlands Ministry of Economic Affairs, Agriculture and Innovation
- Vinod Ralh, IBM
- Chiel Steinvoot, ABN Amro
- Gert-Jan Ter Weeme, Netherlands Ministry of Infrastructure and the Environment
- Gerrit Jan Timmermans, Architecture @ Heart
- Ron Tolido, Capgemini
- Jaap van der Does, ABN Amro
- Dick Van Es, ING
- Adriaan Verdaasdonk, ABN Amro
- Marco Vink, Netherlands Ministry of Infrastructure and the Environment
- Marc Welters, Ernst & Young

The second review produced a number of detailed improvements, and confirmed that the Guide represents a consensus of The Open Group members. The following people made comments or participated in the review teleconferences, which were managed by Raina Wissing of The Open Group.

- Sandra Adamson, Capgemini
- Ram Allampalli, Cognizant
- Leslie Anderson, Raytheon
- Omkhar Arasaratnam, IBM
- Carlos Arevalo, IBM
- Hemant Babtiwale, Capgemini
- Giovanni Ballarini, HP
- Stuart Boardman, Getronics Consulting
- Xicotencatl Bojorges, IBM
- Michael J. Broderick, Capgemini
- Michael Brokmann, IBM
- Wolfgang Bross, HP
- Geoff Burke, British Telecom
- Xijia (Frank) Chen, Cognizant
- Ron de Jong, Oracle
- Prantik Debnath, IBM
- Rakesh V. Dharmala, CA
- Tom Dickinson, IBM

- Steve Dowling, Sirius
- Daniel Downing
- Jonathan Eccles, HP
- Kerry Finn, Raytheon
- Olav Frolich, Devoteam
- Mark Gasser, Boeing
- Shuvanker Ghosh, IBM
- Alan Giblin, HP
- Penelope Gordon, 1Plug
- Ed Harrington, Architecting the Enterprise
- Christine Huang, Deloitte
- Gill Hughes, Capgemini
- Mihai Iancu, CCP Investment Board
- Nilton Ideriha, IBM
- Niek IJzinga, Logica
- Glenn Incorvia, Cognizant
- Pamela K. Isom, IBM
- Heinz Johner, IBM
- Juan Juan, HP
- Radha Kasibhatla, HP
- Venkata R. Kona, Wipro
- Lucky Krishna, IBM
- Joe Lofgren, Raytheon
- Bhavish Madurai, CSC
- Dan Marshall, IBM
- Rafael Araba Moratinos, Oracle
- Tuomas Nurmela, Tieto
- Marlin Pohlman, EMC
- Raffaele Pullo, IBM
- Leonardo Ramirez, Dux Diligens
- Sudhir Rao, Capgemini
- Richard S. Raszka, IBM
- Claude Rioussset, IBM
- Michael P. Rogers, Capgemini
- Rajib Roy, HP
- Tarun Roy, Capgemini
- Nayan Ruparelia, HP
- Ilyas Sener, Capgemini

- Sanjeev Sharma, IBM
- Dushyant Singh, Tata
- Bjoern Steffens, IBM
- Silviu Stoian, Oracle
- Rey Tapia, Ernst & Young
- John Taylor, Logica
- Saji Thoppil, Wipro
- Bugra Uytun, IBM
- Martin Van Den Berg, Sogeti
- Subbarao Varahabhotla, IBM
- Guillermo B. Vasquez, SAP
- Christopher Voigt, IBM
- Wolfgang Von Drews, IBM
- Stuart Warren, Logica
- Bob Weisman, Build the Vision
- Michael Williams, Cisco

Thanks are also due to Annelise Savill of VHP for her encouragement and support, Jan Niessen of VHP for his help with the reviews, Allen Brown and James Scott of The Open Group for their advice and input, and Cathy Fox of The Open Group for her meticulous editorial work.

The Open Group is grateful to all of these people for their contribution to the Guide, and would also like to acknowledge the more than 500 other members of the Cloud Computing Work Group whose activities contributed to the body of knowledge on which the Guide is based, even though they did not play a specific role in its development.

Referenced documents

The following documents are referenced in this Guide:

- [BUSSCEN] Cloud Computing Business Scenario Workshop, Report, August 2009 (R091), published by The Open Group; refer to: www.opengroup.org/bookstore/catalog/r091.htm.
- [BUYERSTREE] Cloud Buyers' Decision Tree, White Paper, July 2010 (W107), published by The Open Group; refer to: www.opengroup.org/bookstore/catalog/w107.htm.
- [BUYERSQ] Cloud Buyers' Requirements Questionnaire, Version 1.0, White Paper, July 2010 (W108), published by The Open Group; refer to: www.opengroup.org/bookstore/catalog/w108.htm.
- [CBUC] Strengthening your Business Case for Using Cloud, White Paper, July 2010 (W106), published by The Open Group; refer to: www.opengroup.org/bookstore/catalog/w106.htm.
- [CHRISTENSEN] The Innovator's Dilemma – When New Technologies Cause Great Firms to Fail, Clayton M. Christensen, Harvard Business School Press, ISBN: 0-87584-585-1.
- [CLOUDROI] Building Return on Investment from Cloud Computing, White Paper, April 2010 (W104), published by The Open Group; refer to: www.opengroup.org/bookstore/catalog/w104.htm.
- [COSO] Committee of Sponsoring Organizations of the Treadway Commission (COSO); refer to: www.coso.org.
- [CSA] Cloud Security Alliance; refer to: www.cloudsecurityalliance.org.
- [D&A] Rethinking Risk Management, Audrey Dorofee and Christopher Alberts, CMU SEI; refer to: www.sei.cmu.edu/library/abstracts/risk/upload/dorofeetutorialIndia09_8819.pdf.
- [DIACAP] US Department of Defense Information Assurance Certification and Accreditation Program (DIACAP), 2007.
- [DMTF] Distributed Management Task Force; refer to: www.dmtf.org.
- [ECLIPSE] Eclipse Foundation; refer to: www.eclipse.org.

- [EUCALYPTUS] Eucalyptus: An Open Source Cloud Computing Infrastructure, Daniel Nurmi, Rich Wolski, Chris Grzegorzczuk, Graziano Obertelli, Sunil Soman, Lamia Youseff, Dmitrii Zagorodnov; refer to: http://iopscience.iop.org/1742-6596/180/1/012051/pdf/jpconf9_180_012051.pdf.
- [EU 95/46/EC] EU Data Protection Directive 95/46/EC, October 1995, European Parliament and the Council of the European Union.
- [EXAMPLES] Downloadable spreadsheets containing data for the large examples in this book; refer to www.opengroup.org/cloudcomputing/doc.tpl?dcat=22&gdid=24635&lastver=Y.
- [FAIR] The Open Group FAIR – ISO/IEC 27005 Cookbook, Technical Guide, November 2010 (C103), published by The Open Group; refer to: www.opengroup.org/bookstore/catalog/c103.htm.
- [FEDRAMP] US Federal Risk and Authorization Management Program (FedRAMP); refer to: www.cio.gov/pages.cfm/page/Federal-Risk-and-Authorization-Management-Program-FedRAMP.
- [FEDSTRAT] Federal Cloud Computing Strategy, Vivek Kundra, US Chief Information Officer, February 2011; refer to: www.cio.gov/documents/Federal-Cloud-Computing-Strategy.pdf.
- [FISMA] US Federal Information Security Management Act (FISMA), 2002.
- [GLBA] US Gramm-Leach-Bliley Act (GLBA), 1999.
- [HIPAA] US Health Insurance Portability & Accountability Act (HIPAA), 1996.
- [ICAEW] ICAEW Risk Management for SMEs; refer to: www.icaew.com/index.cfm/route/120062/icaew_ga/pdf.
- [ISO 31000] ISO 31000:2009, Risk Management – Principles and Guidelines; refer to: www.iso.org.
- [ISAE 3402] International Standards for Assurance Engagements (ISAE) No. 3402.
- [SSAE 16] Statement on Standards for Attestation Engagements (SSAE) No. 16.

- [ITIL] Information Technology Infrastructure Library; refer to: www.itil.org.uk.
- [IYAR] Why Buy the Cow, Subrah S. Iyar, Webex Communications, ISBN: 9780615163130.
- [LONG TAIL] The Long Tail: Why the Future of Business is Selling Less of More, Chris Anderson, New York: Hyperion, ISBN: 9781401309664.
- [MCDONALD] Legal and Quasi-Legal Issues in Cloud Computing Contracts, Steve McDonald, General Counsel at the Rhode Island School of Design; refer to: http://net.educause.edu/section_params/conf/CCW10/issues.pdf.
- [NIST] The NIST Definition of Cloud Computing, Version 15; refer to: www.nist.gov/itl/cloud/upload/cloud-def-v15.pdf.
- [O-ACEML] Open Automated Compliance Expert Markup Language (O-ACEML), Technical Standard, May 2011 (C111), published by The Open Group; refer to: www.opengroup.org/bookstore/catalog/c111.htm.
- [OAuth] An Open Protocol to Allow Secure API Authorization in a Simple and Standard Method from Desktop and Web Applications; refer to: <http://oauth.net/>.
- [OGF] Open Grid Forum; refer to: www.ogf.org.
- [ORANGE BOOK] The Orange Book: Management of Risk – Principles and Concepts; refer to: http://hm-treasury.gov.uk/d/orange_book.pdf.
- [PCI-DSS] Payment Card Industry Data Security Standard (PCI DSS); refer to: www.pcisecuritystandards.org.
- [RISK] Risk Taxonomy, Technical Standard, January 2009 (C081), published by The Open Group; refer to: www.opengroup.org/bookstore/catalog/c081.htm.
- [SEI RISK] CMU SEI Risk and Opportunity Management; refer to: www.sei.cmu.edu/risk/.
- [SNIA] Storage Networking Industry Association; refer to: www.snia.org.
- [SOX] US Sarbanes-Oxley Act (SOX), 2002.

- [SWITCH] The Big Switch: Rewiring the World, from Edison to Google, Nicholas Carr, ISBN-10: 0393062287.
- [TOGAF] The Open Group Architecture Framework (TOGAF); refer to: www.opengroup.org/togaf.
- [TPC] Transaction Processing Performance Council; refer to: www.tpc.org.
- [UNIX] The UNIX System; refer to: www.opengroup.org/unix.
- [UPTIME] Uptime Institute; refer to: www.uptimeinstitute.org.

Chapter 1

What is cloud?

Cloud computing is the latest major evolution in computing. It is a paradigm where computing resources are available when needed, and you pay for their use in much the same way as for household utilities. Just as water is piped to your home and you pay for as much or as little as you use, cloud computing resources are available whenever needed and charges are based on how much you use them. When you turn it off, the water that you would have used is available for use by others and, in the same way, shared cloud resources can be used by others when not used by you.

Widespread cloud computing is made possible by the Internet, and this is the most common way of accessing cloud resources. Intranets and dedicated networks are sometimes used too, in the case of a private cloud, for example.

In an enterprise that has complex and expensive IT systems to support its business processes, who would not be attracted by the idea of just being able to pay on-demand for someone else to provide IT services without being concerned with the details of how it is done? Who would not welcome having several potential IT service suppliers, giving a competitive choice? Perhaps you no longer need an internal IT department, a business within your business of people who aren't core to your products and services. You can stop worrying about hiring and retaining a workforce with IT skills that are in short supply and therefore at a premium. And you can escape from the confines of application upgrades and hardware obsolescence.

Put like this, it all sounds so simple. But what does 'cloud computing' really mean?

This chapter:

- Takes a look at the leading consensus definition of cloud computing, from the US National Institute of Standards and Technology (NIST), and explains its concepts

- Summarizes the key differences between cloud computing and traditional IT
- Describes how cloud computing is used, and how it impacts providers and consumers of cloud services and of added services implemented using the cloud
- Describes the impact of cloud computing on business

It will give you an insight into the potential for cloud computing to transform business. The next chapter (Chapter 2) explains the different reasons why enterprises are using the cloud, or providing cloud services. The subsequent chapters will help you to understand more specifically how cloud computing can benefit your business.

1.1 NIST definition of cloud computing

The definition of cloud computing provided by NIST has gained significant traction within the IT industry. According to this definition (see [NIST]):

“Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics, three service models, and four deployment models.”

A note to the definition says that:

“Cloud computing is still an evolving paradigm. Its definitions, use-cases, underlying technologies, issues, risks, and benefits will be refined in a spirited debate by the public and private sectors. These definitions, attributes, and characteristics will evolve and change over time.”

This is undoubtedly true. Nevertheless, at this point in time, the NIST definition is widely accepted and is increasingly regarded as authoritative.

Cloud computing is often confused with other similar computing paradigms. Examples of common computing models that are not cloud computing are: peer-to-peer networks, Service-Oriented Architecture (SOA), grid

computing (for example, the SETI@home project), and network asset sharing (networked printers, NAS disk drives, etc.).

The basic cloud computing model is shown below (in Figure 1.1). Servers, storage, applications, and services are accessed via a common network. They are shared between organizations, and accessed by users or applications. The users may be members of the organizations working on-the premises, remote workers, customers, or members of the general public.




Figure 1.1: Basic cloud computing model

Cloud computing resources can be rapidly provisioned and released with minimal management effort or service provider interaction. This means that an organization can use more or fewer servers, stores, applications, or services, and can configure the ones it uses to meet its requirements, as and when it wishes to do so, and without major effort.

Cloud computing has five essential characteristics. They are on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service. These are the features that distinguish it from other computing models.

It has three service models. These are the basic kinds of service that cloud service providers provide. They are Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS).

It has four deployment models, corresponding to different ways in which cloud resources are owned and operated. They are private cloud, public cloud, community cloud, and hybrid cloud.

1.2 Essential characteristics

The five essential characteristics of the NIST definition are shown below (in Figure 1.2) and described in the following sections. The definitions of the characteristics in these sections are all taken from the NIST definition.




Figure 1.2: The essential characteristics of cloud computing

1.2.1 On-demand self-service

“A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service’s provider.”

Consumers must be able to obtain cloud services – at the infrastructure, platform, or application level – whenever they want, without requiring significant assistance.

Often, these services must be paid for. The buyer can typically set up an account with the seller, establish security and billing credentials, and then select and schedule the use of the cloud computing resources on sale. This is generally done using an easily accessible and user-friendly online system. In the case of IaaS, for example, this might enable the user to start virtual machines, assign network addresses, and allocate storage.

In some situations, particularly with private cloud, payment is not needed. The service provider will still give the consumer the ability to configure resources and schedule their use.

For the provider, on-demand self-service requires that procurement, account management, service instantiation, security control, service management, metering, billing and payment mechanisms are established. These mechanisms will interface with operational systems so that services are created, started, run, and stopped in accordance with the consumer's instructions. Use of a service by a consumer might last only for minutes, or for weeks, months, or years.

Although the concept of cloud computing gives the illusion of infinite resource, clearly this in reality does not exist, and most certainly does not exist for all consumers at the same point in time. How well providers are able to forecast demand, and how efficient they are at provisioning their services, are fundamental to their ability to meet consumer expectations and service levels.

Providers typically set usage limits in line with their ability to provision resources and with a buyer's credit-worthiness. They may determine that usage requests that require a large amount of application or infrastructure resource, or that cost more than a set amount, are subject to minimum notice, pre-reservations, or additional financial vetting. Various means may be used to help preserve service levels, such as offering reduced prices for advance purchases or for buyers who are willing to endure service interruptions or reduced performance in the event of resource constraint.

1.2.2 Broad network access

“Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and personal digital assistants (PDAs)).”

Network access is needed to establish the initial provider/consumer relationship, for subsequent use of the cloud services themselves, and for use of added services that the consumer may implement using the cloud services. For example, a company might use cloud services to implement a web site to give its customers product information. That company needs network access

to purchase and configure the cloud services, and to manage and operate its web site, and its customers need network access to use the web site.

The user of a cloud service or added service might have a PC or a device such as a tablet, a PDA, or a mobile phone. These could have browsers or be browser-less devices. Or the cloud services could be integrated into a consumer enterprise's IT architecture, with access from large and sophisticated computer systems.

The cloud services must be accessible through standardized mechanisms. This introduces an important consideration for all cloud implementations – the use of standards. Without adherence to standards throughout the technology stack, from the network level up to the client access and presentation level, accessibility from such a variety of devices and applications would inevitably be reduced and broad network access would not be achievable. Use of standards is fundamental to this.

Commercial companies generally want maximum market access. The greater the accessibility of their services, the greater is the potential for sales. The broad network access characteristic of cloud computing means that a company can implement added services that can be successfully used by anyone, anywhere on the globe, using a variety of devices.

1.2.3 Resource pooling

“The provider’s computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location-independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or data center). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.”

A traditional IT model is generally centered on the provision of IT services for one enterprise. Because they can assign pooled resources dynamically to meet demand, cloud service providers can maintain maximum service levels with minimum resources. For consumers, this means high QoS at low cost. It is a major reason why cloud computing, unlike other initiatives, is expected to succeed.

Having consumers from different industry segments and countries can help significantly. Industries and countries have varying usage patterns because of climate and cultural differences. When resources are shared, the patterns are combined, with usage peaks in one industry or country coming at the same time as periods of lower use in others.

Computing resources may be shared at the infrastructure, platform, or application level. Most cloud computing users share infrastructure and possibly platforms. That said, sharing platforms and applications does not necessarily mean users sharing instances of operating systems and applications programs. From an application standpoint, enabling sharing could require extensive and fundamental re-architecting of the software. This may not be worth doing, given that the key benefit of efficient resource use can be obtained just by pooling infrastructure and platform resources.

Efficient resource use is one side of the resource-pooling coin; multi-tenancy is the other.

Multi-tenancy means that a single instance of a computing resource serves multiple client organizations (the tenants) providing a separate environment for each. Examples of resources include instances of infrastructure, platform, software, and application.

Multi-tenancy tenant isolation often relates to fault isolation, resource isolation, and security isolation. These capabilities enable tenants to have secure, available environments regardless of other tenant behavior.

How multi-tenancy is enabled depends on the service model. In the case of IaaS, multi-tenancy of the infrastructure is enabled by the virtualization of the infrastructure resources. For PaaS, multi-tenancy of a platform is enabled by the platform software providing separate environments for its user organizations. In this sense, a multi-user operating system can be regarded as multi-tenant by definition. For SaaS, multi-tenancy of a software application depends on the application being designed to partition its configuration and data for the client organizations. Most of today's applications are designed for a single tenant, and cannot be changed to multi-tenant operation without significant re-architecting.

From the user's point of view, it can be hard to distinguish a multi-tenant application from a single-tenant application running on a multi-tenant platform, or to distinguish a multi-tenant platform from a single-tenant platform running on multi-tenant infrastructure.

Each consumer must consider whether there is any impact from sharing with other, unknown, consumers. It may be hard to establish the degree to which sharing takes place, and impossible to discover who the other consumers are. Providers give varying amounts of information about their systems, which may include very specific product configurations and bespoke enhancements to standard products. They may offer guarantees of service and security levels, or provide very little by way of assurance. You must assess the risks and make your own judgment.

The concept of resource pooling includes location independence. For providers, having resources in convenient locations means lower costs. Having them in diverse locations means that services can be maintained in the event of loss of a data center, of power, or of network connectivity. Again, this translates to higher quality and lower cost for consumers.

As far as their use of the services is concerned, consumers have no need to know what the underlying resources are, or where they are located. But there are other reasons – such as data security, compliance with regulations, and performance in accessing data – why location may become highly relevant. Because of this, providers may give a choice of location, at least to the extent of continent or country, or may say where they operate.

For SaaS there may be certain levels of protection effectively built into the application service, by routine back-ups and multi-data center and even multi-country hosting. For PaaS and IaaS, while the possibility to construct such resilience may exist, the consumer will probably have to select and configure the individual components required, assign geographically dispersed servers to construct failover configurations, and copy the data, much as for in-house IT solutions.

How a provider's resources are distributed geographically, the amount of protection against various disaster scenarios that this provides, and the legal issues associated with keeping data in the countries concerned, will

be matters of detail for the consumer to investigate. See the discussion of Establishing Requirements (in Section 4.2) for more on the legal issues.

1.2.4 Rapid elasticity

“Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.”

One of the key benefits of cloud computing is the ability to have a flexible computing service which can expand or contract in line with business demand, giving you capacity which would be impossible to generate from an in-house implementation without significant investment in resources.

Elasticity has three major features:

- **Linear scaling**
The service can scale, independent of the number of users or workload size (subject to operating condition limits). The performance experience for one of a thousand users is the same as for a single user. This is not the case for most traditional systems, which have non-linear scalability rather than elasticity.
- **On-demand utilization**
Allocation of virtual resources follows the demand profile exactly, so that the user appears to have 100% utilization of the service.
- **Pay-as-you-go**
Payment for IT resources is on a per-use basis using an OPEX style charging principle. This means that the asset ownership is with the service provider, and the user pays for consumption of the service on the basis of the resource units consumed. Typically, a cloud service has logical units of resource defined by the service provider based on a standard configuration.

Resource pooling helps providers to achieve elasticity. A resource that is no longer needed by one consumer can be allocated to another consumer that needs more resources.

Consumers must understand how elasticity is provided and how different usage levels are priced, in order to exploit this characteristic to best advantage. Pricing can be complex, and can vary widely between suppliers. Use of applications that provide the same function may be charged by transaction, by number of users, by data quantity, or by infrastructure consumed. Pricing of infrastructure resources may take into account factors such as virtual machine size, memory and disc size, and network usage.

For SaaS, this may be relatively simple, with charging based upon real transaction or user utilization. The consumer does not need to worry about resource configuration, and probably has no ability to change it anyway.

With PaaS or IaaS, the consumer may need to design an effective application and infrastructure solution to exploit rapid elasticity. For example, IaaS providers let you schedule and run as many virtual servers as you wish; it is for you to determine how many are required to host your application. You may need to perform your own monitoring of capacity and utilization to ensure that excess virtual servers are stopped when not required and additional ones are started when demand rises.

For organizations selling cloud computing services, the challenge is to attract and retain buyers, and to do so while making a healthy profit. Forecasting how many buyers will buy how much and when is analogous to customer-demand forecasting by utility companies; it is radically different from the same exercise undertaken within a single organization with established patterns of business activity. Over-capacity will cost in respect to infrastructure and applications deployed but not sold; under-capacity may lose customers and thereby revenue. Sellers can mitigate some of the issues by using tactics such as buyer behavior analysis and flexible purchasing arrangements with hardware and software vendors.

1.2.5 Measured service

“Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported providing transparency for both the provider and consumer of the utilized service.”

If services are to be charged on a per-use basis, it is immediately obvious that usage must somehow be measured. Consumers require sufficient measurements from their cloud computing service providers to enable them to make purchasing and operational judgments. First, they want to select the optimum service to meet their needs. Then, they want to know whether it is performing in accordance with their SLAs, and whether their use of it is as expected.

Sellers of cloud services must provide sufficient information about their charging regimes up-front to allow purchasers to make informed choices. On an ongoing basis, they must provide accurate accounting information to support their bills, and give sufficient usage information to allow solutions to be managed operationally.

Consumers may wish to use this information in systems management and financial accounting applications. They may wish to aggregate or sub-divide it if they in turn provide services to other organizations. The information must be in an appropriate form to allow for this.

Usage of different components may need to be measured separately. In IaaS, charges are often calculated for storage occupied, network data transfers, IP addresses, virtual servers, etc. Each of these components will require measurement to ascertain number, size, quantity, and usage over time as appropriate to the charging regimes in place. Similar measures will be required for PaaS. Measures for SaaS might include per-user and per-software-function values, which might in turn be calculated from measures of infrastructure occupancy.

1.3 Service models

NIST defines three service models. They parallel the layers of a traditional computing environment. They are illustrated in the figure below (Figure 1.3) and described in the following sections.




Figure 1.3: Cloud service models

1.3.1 Cloud infrastructure as a service (IaaS)

“The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).”

Instead of owning, managing, and operating your own computer hardware, you can use computers that someone else owns, manages, and operates.

Computer leasing and outsourcing became popular back in the early days of computing with companies that wanted to use computers but didn't want to bear the capital cost of owning them or to have the hassle of operating them. Since then, two technical developments have made possible IaaS as we have it today, providing the same benefits but in radically improved form: the Internet and virtualization.

The availability of high-bandwidth data communications over the Internet makes it almost as easy to use a computer in a different continent as one on your own premises. This eliminates any problems with having operations staff from other companies onsite, and means that the company running your computers can organize itself efficiently on its own premises, taking advantage of locations where staff costs and operating costs are low.

Virtualization is a technique by which what appears to be a single computer or storage module is in fact not implemented as a single physical resource but uses part of the capacity of one or more physical resources. When you use an IaaS ‘machine’ or ‘storage block’, it is very unlikely that you will have a dedicated computer or disk drive.

There are a number of virtualization techniques and products that enable this. These have been developed as a result of commercial research and academic studies at leading institutions including the Cambridge University Computer Laboratory in the UK, the Center for Research on Computation and Society School for Engineering and Applied Sciences at Harvard University, the University of California Santa Barbara, the University of Berkley RAD Laboratory, and the Massachusetts Institute of Technology (MIT).

Generally, in these products, a hypervisor (consisting of a program possibly supported by special-purpose virtualization hardware) presents virtual hardware resources to the guest operating systems, allowing industry standard servers and their attached network and storage to function as unified resource pools. Thus, a number of physical resources are collected together and, through a single hypervisor interface, are presented as a collection of virtual machines, on which the guest operating systems run. The hypervisor also monitors the execution of the virtual machines on the physical resources, and provides configuration and management capabilities. An example is the Eucalyptus open source cloud computing infrastructure [EUCALYPTUS].

It is these virtualization techniques that make effective resource pooling and elasticity possible for IaaS.

1.3.2 Cloud platform as a service (PaaS)

“The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.”

You can pass to a service provider the burden of owning, managing, and operating systems software, as well as the hardware itself.

The kind of ‘as a service’ determines the division of responsibility between consumer and provider. The PaaS definition implies that the service provider is responsible for run-time monitoring and management. (The consumer defines configurations, but it is then up to the provider to follow them.) This requires middleware. Furthermore, there has to be a deployment mechanism for applications as part of PaaS. PaaS providers typically have Linux or Windows operating systems for their customers to use, often with SQL databases and support for standard programming languages such as Java. They can also provide middleware and web-hosting platforms with servlet containers and support for server-side scripting languages such as PHP.

1.3.3 Cloud software as a service (SaaS)

“The capability provided to the consumer is to use the provider’s applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.”

In addition to applications accessed by browsers, there are cloud services accessed by browser-less devices. Cloud applications for mobile devices constitute a growing and popular market.

Ultimately, you can let the service providers run all of your software. You may need different applications from different vendors. In that case, you may have to worry about integration.

An increasing range of applications is available on the cloud, from storage and exchange of personal information on social networking sites, to office applications, and business applications such as product quality management.

1.4 Deployment models

NIST defines four deployment models: private cloud, public cloud, community cloud, and hybrid cloud.

1.4.1 Private cloud

“The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premise or off-premise.”

A private cloud, where the infrastructure is operated and managed by a third party, is sometimes called a virtual private cloud, particularly where the third party is a public cloud provider that dedicates a part of its cloud infrastructure to public use and part to private use.

Private cloud gives an organization the benefits of cloud computing, without the restrictions of network bandwidth, security exposures, and legal issues that using external resources might entail. It can also have better security, accountability, and resilience than public cloud, because use can be controlled and managed.

The cost of private cloud can be borne by the organization as a whole, rather than by individual users. Not having to pay or account for resources consumed can make using the cloud simpler and easier.

An ‘organization’ can represent anything from an individual to a large commercial corporation or government department. Private cloud can be especially useful for larger enterprises, because they can gain greater economies from shared use of pooled resources. Also, at least at present, the cost of designing and implementing a private cloud may be prohibitive for a small organization. (This is likely to change as more vendors offer private cloud as off-the-shelf products.)

The advantages of private cloud should be balanced against possible disadvantages. Some things to consider are:

- **Capital investment:** Significant capital investment is required for the parent organization. Unless there is senior sponsorship in the company and a robust business case (or an appetite for upfront investment), this can be enough to stall the activity.

- **Time-to-market:** Most companies will take 6-36 months to establish a usable private cloud. Those that take less time usually forget key requirements such as auto-provisioning, governance, chargeback, monitoring, and adequate service operations. Significant time can be added if high availability is required, as anything above 99.9% usually requires creation of more than one location for the data center.
- **Technology:** An organization needs to make important decisions such as buying best-of-breed point technology solutions, and this inevitably means significant integration work or buying a single-vendor solution, which might not be best fit in all areas.
- **Learning curve:** Cloud vendors have learned some very valuable lessons, which even strong technology companies will not be able to bake into their solutions. They are also very unwilling to share best practices with others.

1.4.2 Public cloud

“The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.”

‘Public’ does not mean ‘free’. Public cloud providers may offer some services free of charge, but in general they charge enough on average to at least cover their costs.

Also, ‘public’ does not mean that user data is visible to the public at large. Cloud providers implement security mechanisms to control access to data.

The main benefit of using a public cloud, as opposed to creating a private cloud, is easy and inexpensive set-up. The provider has done the work needed to create the cloud; the consumer just needs to do an additional amount to configure the resources to be used. The consumer also benefits from the economies of sharing resources with other consumers (though multi-tenancy may have a downside too).

1.4.3 Community cloud

“The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on-premise or off-premise.”

Community cloud, as defined here, has similarities to both private and public cloud. Like private cloud, it can avoid network bandwidth, security exposures, and legal issues that arise from using external resources, and its use can be controlled and managed. Like public cloud, it makes set-up easy for individual organizations, and it provides more efficient use of pooled resources for the whole community than any of its members could achieve individually.

While the burden of creating and managing the cloud is lifted from the shoulders of each member organization, this has to be done by the community as a whole. Charging mechanisms and governance procedures must be established to enable this.

Case 24: Virtual Learning in Cloud Computing in Use (Appendix A) provides an example of community cloud.

Note that social networking and similar services do not really satisfy this definition. Here, the infrastructure is owned by a single organization and shared by a community of end users. This is classed as public cloud.

1.4.4 Hybrid cloud

“The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).”

A hybrid cloud may be coordinated by a broker that federates data, identity, security, and other details.

Another scenario is that an enterprise has a private cloud but also uses a public cloud, perhaps for particular applications, or perhaps as a back-up or to handle peaks of load. In this model users typically host non-business-critical information and processing in the public cloud, while keeping business-critical services and data in their control in the private part of the hybrid.

1.5 Extending the NIST model to business processes

The phrase ‘as a Service’ has become something of a buzzword. We hear of such things as Storage as a Service, Information as a Service, and Integration as a Service. In many cases, these are sub-categories of the three NIST service models, IaaS, PaaS, and SaaS.

There is, however, one area in which the cloud community is talking about a major extension: the provision of cloud services that carry out business processes, such as payroll, CRM, billing, HR, order taking, and information delivery. This is a new service model: Business Process-as-a-Service (BPaaS). It is differentiated from SaaS because it includes services partly performed by people, not just by applications software.

In this model, the consumer has the ability to use the provider-defined business processes running on cloud services. The business processes interact with various client devices through lightweight interfaces such as a web browser or email. The consumer does not manage or control the underlying cloud platform and infrastructure, including network, servers, operating systems, storage, and BPM platform. Nor does the consumer manage or control individual business processes and underlying application capabilities, with the possible exception of limited consumer-specific process configuration settings.

While BPaaS is potentially a very significant development, this Guide focuses on the original cloud service models of IaaS, PaaS, and SaaS.

1.6 Comparison with traditional environments

The key practical differences between traditional computing environments and cloud computing are shown below (in Table 1.1).

Characteristic	Cloud Computing	Traditional	Comments
Time before service can be accessed	Minutes/ Hours	Days/Weeks	Once the cloud computing environment is set up initially, you can gain access faster than in traditional environments where lead time is needed for installation, set-up, and configuration.
Capital Expenditure (CAPEX)	Pay-as-you-go, Variable	Upfront cost, Fixed	The pay-as-you-go model for cloud computing reduces or eliminates the large upfront costs incurred in procuring hardware and software and standing up traditional environments.
Economies of scale	Yes, for all organizations	For large organizations only	Cloud computing not only provides cost advantages in procurement of hardware and software, it also provides cost advantages from improved productivity. Traditionally, lessons learned from one environment must be duplicated in other environments but, with cloud computing, once the best practices are applied they benefit all consumers.
Multi-tenancy	Yes	Generally no, but can be found in application hosting	Multi-tenancy properly applied to cloud computing services allows providers to host multiple consumers effectively across shared resources. While it is more readily enabled in IaaS through the use of virtualization, PaaS and SaaS providers may need to undertake significant re-architecting of their platforms or applications to apply multi-tenancy to these elements as well as to infrastructure. Where this has not been undertaken, consumers may find that their platforms and applications are not as elastic or cost-effective as anticipated.
Scalability	Elastic and Automatic	Manual	Cloud computing resources can often be scaled up or down automatically, whereas human intervention is usually needed to add hardware and software in traditional environments.
Virtualized	Usually	Sometimes	Cloud computing environments are usually virtualized, whereas traditional environments include a mix of physical and virtualized infrastructure.

Table 1.1: Practical differences between cloud computing and traditional environments

1.7 Providing and using cloud services

Cloud computing is about the provision of services. A service has a provider and a consumer. It exposes capabilities that the provider has that are of value to the consumer. For cloud services, these capabilities are those of:

- Owning and operating computer hardware (IaaS)
- Owning and operating computer hardware and systems software (PaaS)
- Owning and operating computer hardware, systems software, and software applications (SaaS)

In many cases, an organization uses cloud services in order to provide added services to the people and organizations that it deals with. This is obviously the case when a company buys IaaS or PaaS from a cloud provider and runs a software application on top of it. The cloud service consumer develops the added service and provides it for the benefit of end users. It is also the case when a company buys SaaS from a cloud provider and uses that service to support its business operations. And it is the case when an enterprise's IT department uses the cloud to provide the services needed by other departments. Using the cloud to provide added services means that those services must be developed on, or ported to, the cloud.

In other cases, for example, when a company uses an SaaS CRM solution, organizations and individuals simply use cloud services and added services to support their business and leisure activities. This has many benefits, but there are some pitfalls to avoid.

1.7.1 Providing cloud services

The provider of a cloud service has control over a set of resources, and makes them available to consumers of the service, in a way that has the essential characteristics of cloud computing.

The resources may include processors, data stores, system programs, application programs, and networks. They may be owned or leased by the provider and be on the provider's premises, or the provider may have control of them through a contract with another cloud service provider.

The resources are made available to consumers under a contract. For large-scale use, this may take the form of a signed legal document. In other cases, the consumer typically checks a box on a web form. The provider agrees to

provide the service, under certain conditions. The consumer agrees to accept the conditions when using the service and, in many cases, to pay for the service.

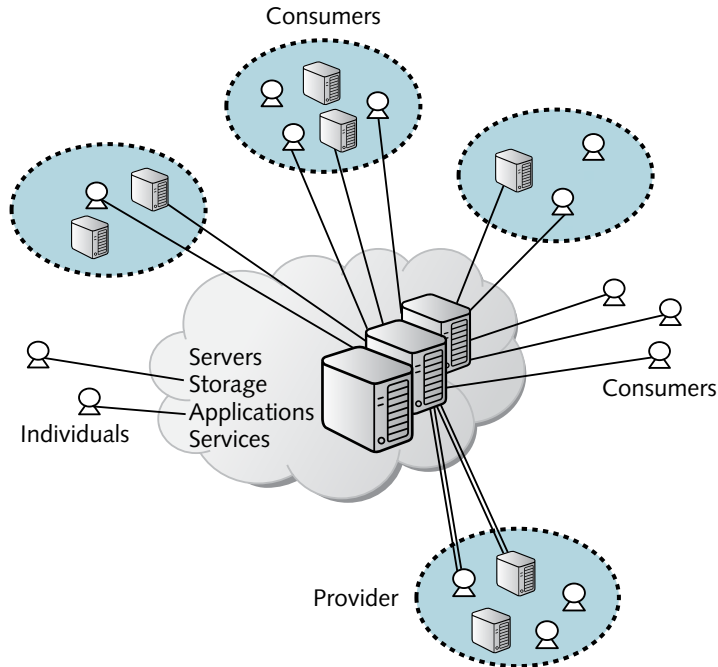


Figure 1.4: Providing cloud services

On-demand self-service requires something at the point of interaction between the consumer of a cloud computing service and the provider to enable the service to be contracted for and instantiated with minimal effort and interaction between these two parties. The typical means for this is a web site that enables the consumer to:

- Sign up to the cloud service
- Access billing and payment information and mechanisms
- Access administrative capabilities, particularly for security and identity management
- Operate the cloud service – adding, running, stopping, modifying, and removing resources as required

This web site is used to establish a provider/consumer relationship and contract. It is not used for the subsequent consumption of added services by end users. For example, if a company runs a web site using resources from a

cloud service provider, visitors to the web site may have no idea that it was provided through use of cloud computing infrastructure, and would not interact directly with the cloud service provider.

Beyond the essential capabilities to allow a cloud computing resource to be purchased and run, additional capabilities may be present to provide other functions common within IT operation, such as monitoring, reporting, SLA management, error reporting, and bug fixing.

1.7.2 Providing added services

The consumer of a cloud service may use it to provide an added service, as shown below (in Figure 1.5). The added service is used by the customers of the provider of the added service, not by the provider itself.




Figure 1.5: Providing added services

The provider of added services contracts with cloud providers to obtain use of the resources that they control. The added services may in turn be cloud services. For example, an SaaS provider might use resources from a PaaS provider. Generally, the added services are software services, but they do not

necessarily have the essential cloud computing characteristics. For example, they may not be available through on-demand self-service, or they may not be elastic.

Many companies, for example, use cloud services to implement web sites, and deliver added services such as product ordering through those web sites. The end users use the web sites in the normal way; there is no possibility of an end user increasing or decreasing web site capacity (as would be possible with a cloud service), or any need for this. For these added service providers, the cloud is simply a convenient way of obtaining computing resources to support their business activities.

1.7.3 Developing added services

Before an added service can be provided, it must be developed, or ported to the cloud from in-house systems or elsewhere. For the developer, the use of cloud services will require care:

- Thought and effort are needed to make best use of its characteristics to obtain maximum benefit.
- The contractual terms and their implications for non-functional requirements must be well understood.
- It is often necessary to integrate the added services that use the cloud with existing traditional in-house IT services.
- There are differences between the development facilities used in a cloud environment and those of a typical in-house IT department.
- There are areas where the developer is dependent on the cloud provider, and must allow for provider actions that he or she cannot control.

Developing for consistent performance will require different architectural approaches to those typically employed for in-house IT environments, where fixed resources are often provisioned for individual applications and services. Developed solutions may need to interact with the provider's scheduling and resource provisioning services to add or remove resources. The means of engaging with different providers may be radically different.

If services are to be able to respond to demand and meet their requirements for availability, then the demand must be monitored and resources scheduled to satisfy it. Developers should be looking to exploit cloud computing characteristics to manage volatility, scheduling resources in response to

demand to maintain service levels, but avoiding over-capacity. For service availability, applications may be hosted upon multiple clustered servers, and mechanisms must be provided for back-up and recovery of data. Facilities may need to be developed for monitoring, alerting, and scheduling of maintenance operations. Developers will rely upon the presence of service metrics to enable effective management of resources.

Also, the developer must cater for the different pricing models of different providers. For one model, it may be necessary to avoid excessive interaction between servers because these entail high network charges; for another, this may not be a problem, but there could be high costs associated with end-user interaction.

An in-house IT department usually has established procedures for providing help, fault reporting, bug fixing, setting up development environments, testing, and cutting over to live service. Completely different methods, tools, and procedures will in all likelihood be required when developing upon a cloud computing service, whether this is IaaS, PaaS, or SaaS. The whole service development and management lifecycle may be impacted.

Effective development requires that developers understand and exploit the platform upon which they are developing. As with any new development platform, they will need time to learn how to use cloud development facilities. Support by cloud providers for standard development platforms such as Eclipse (see [ECLIPSE]) helps to overcome this problem.

The nature of cloud will be more suited to the exploitation of offshore development resources which many companies are leaning towards, as these are generally geared for a global marketplace.

Developing at a distance may have unforeseen impacts upon the ability to load and extract test data and to migrate live data from existing systems. This may be a concern if an application needs to store or retrieve significant quantities of data and these are dispersed around the globe.

When what you use is metered and you pay for it, the amount of use becomes a major concern. Developers' use of cloud resources may need to be monitored. For example, storing excess copies of data in an on-premise

development is merely a maintenance and management inconvenience, but in a cloud computing development this will cost money.

Developers working in an in-house IT environment are often able to dictate product choices, including versions and configuration parameters. In a cloud computing environment, many of these choices may be restricted by the provider. Changes to the environment such as product version upgrades are likely to take place at the provider's discretion, irrespective of the lifecycle stage of the developer. This could mean significant re-work or re-testing. With IaaS, developers may still have significant control over product and platform choices, but they will have little or no control with PaaS or SaaS.

Fault reporting, bug fixing, and change control for a cloud service will be carried out by the provider. Developers depend on the provider performing these processes effectively.

The relationship of a developer to the cloud computing service provider is different from that of a developer to an internal IT department service provider. The developer has much less ability to influence the provider, but is still dependent on the provider in many ways; for example, in support of fixing problems.

Testing, including functional, performance, and security testing, will be more challenging with cloud service providers, given the relatively closed environment that they generally provide. With SaaS, some testing may need to be done in the production environment.

Quality of the development environment is an important consideration. It should be taken into account when procuring cloud services. In some cases, it may be more important than the quality and cost of the cloud services themselves. A cheap and powerful deployment environment may not make up for increased application programming costs.

1.7.4 Using cloud services

An end-user organization, its customers, and other people that interact with it see the cloud software services, rather than the servers or storage. Those services may be provided by the cloud provider, or by added service

providers. There may be several clouds, several cloud providers, and several added service providers.




Figure 1.6: Using cloud services

If you are a consumer, the ability to connect to and use a service with a device of your choice, from a location of your choice, and using a network connection of your choice may be a significant factor in choosing your provider. The greater the limitations placed upon service access and consumption, such as the use of a particular web browser or web browser version, the more difficult, restrictive, and therefore less valuable the service.

Users of multiple services from different providers may find themselves the victims of conflicting demands. For example, if two services require different browser clients, the user may not be able to accommodate both.

For enterprises with established in-house IT, integrating cloud services with existing systems and service management tools is important, and may not be easy. Enhancing management systems to accommodate cloud services may

require significant work. The lack of standards in this area will be an ongoing obstacle for integration of service management functions. Facilities that are routinely taken as given in on-premise developments, including major non-functional facilities such as authentication and single sign-on, may need enhancement to work with cloud services.

Although managing the complexity of client access software and backend application versions is nothing new to those used to the traditional IT environment of enterprises today, dependence on services from a cloud service provider can bring new difficulties. Established software management and version control policies will no longer apply. Governance practices will also have to adapt. It will no longer be possible to schedule upgrades and changes when required, avoiding key business processing events. Understanding providers' policies for change notification and customer service is critical.

Having different providers for different applications can be a problem, however. It is too easy for them to be 'information silos', so that you do not have integrated access to information across your enterprise. SaaS provides many benefits to its customers, but also provides challenges. These could be major commercial opportunities for systems integrators – and corresponding costs for user enterprises.

There could be integration difficulties at infrastructure and platform levels as well as the application level. Like quality of the development environment, the ability to integrate a cloud provider's services with those of other cloud providers and with those of in-house IT departments is an important consideration that should be taken into account when procuring cloud services.

1.8 The impact of cloud computing

Most people who have water on tap take it for granted. Few have ever had to go to a well or communal standpipe for their water, or know what it is like to have to do this every day. Having water on tap brings a higher quality of life. Freed from the need to devote time and effort to basic but necessary tasks, people can do things that they want to do, that make their life better. In the same way, cloud computing frees enterprises from the need to devote time

and effort to operating and managing their IT, and enables them to focus on their core activities.

The knock-on effects on the wider economy are huge, and very beneficial. As businesses become more effective at what they do, they have capacity to do more. The whole economy grows. Municipal water is an essential pre-requisite for a civilized society today. In the same way, cloud computing will be essential for the civilized societies of the future.

1.8.1 New business paradigm

The evolution of business and IT assets and operations is a continuous process. Today, many factors are involved beyond the disruptions caused by technology developments.

Economic, environmental, and global activities shape regional markets, products, and services in many industry sectors. Government legislation and investment priorities drive standards and commercial behaviors. The Internet, mass media, and collaboration create new access channels that enable market developments hitherto constrained by geography and performance limitations.

Cloud computing provides a new business paradigm for resources. It enables organizations to create and use IT and business services on-demand from optimal sources to maximize utilization and cost-effectiveness. This can be between enterprises or within a single enterprise.

In a business environment, the business processes are how work gets done. They are supported by applications that manage information content and perform transactions. These are in turn supported by a platform and infrastructure that provide storage, processing, and communications. Cloud computing enables businesses to create and use services on-demand, through cloud SaaS, PaaS, and IaaS. Examples of such services include business services, application software services, integration and development services, and infrastructure services. The applications, platform, and infrastructure thus can be replaced by cloud services, as shown below (in Figure 1.7).




Figure 1.7: Use of cloud services to support business processes

The cloud services that support the business processes can be inside the organization, in an internal private cloud, or outside it, in a public cloud or virtual private cloud, as illustrated in the figure below (Figure 1.8). Or they can be in a hybrid cloud, a combination of public and private clouds.

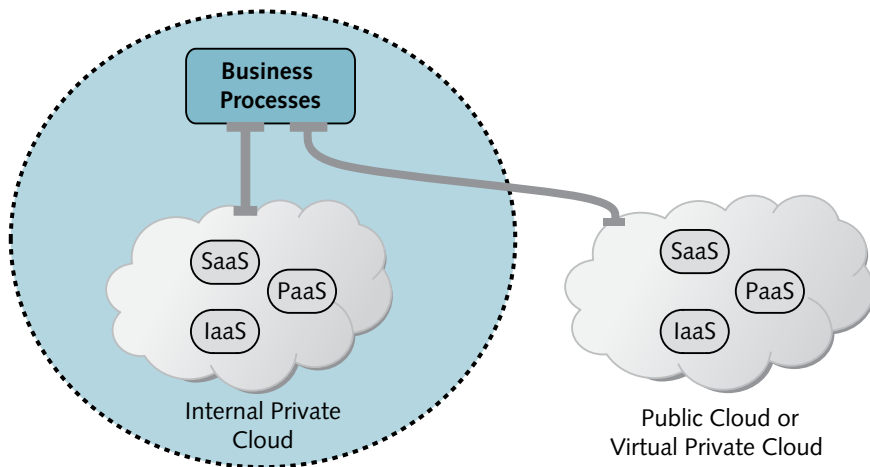


Figure 1.8: Internal and external cloud services

Cloud computing enables businesses to think and act beyond the ‘four walls’ of the company through exchange of services. They can access marketplace best practice solutions, and select effective IT services from multiple sources to meet their needs faster and at lower cost.

1.8.2 Cloud ecosystems

An ecosystem exists where participants in a defined market have integrated business processes and use common standards for exchange of information, products, and services.

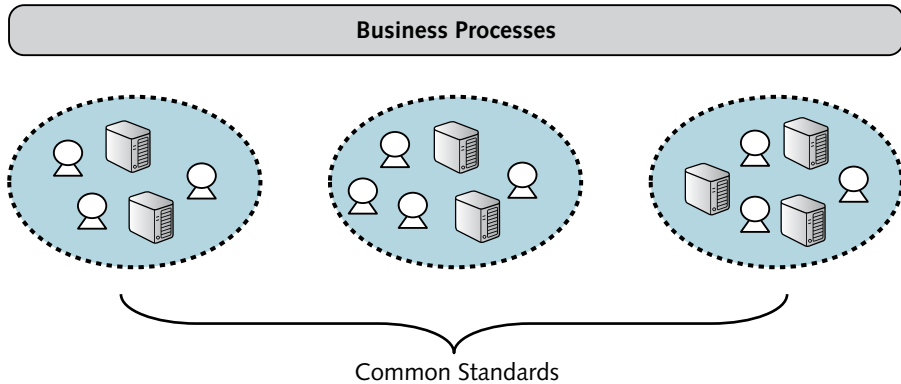


Figure 1.9: Business ecosystem

In today's world, companies are participating in highly collaborative ecosystems providing their specific expertise to create end-to-end services. This will become more important in the future.

SOA and Web 2.0 were milestone developments in the IT industry, while Business Process Management (BPM) has been a major step toward standardized business services automation. With cloud computing, standards and technological developments come together to create an environment in which integrated business processes are supported by software services performed within and between enterprises.

In such an ecosystem, providers and consumers of cloud services participate in common business processes. We have become used to the cloud as a means of sharing information; it is now commonplace, for example, for people to share photographs using social network websites. In cloud-based business systems, this is taken to a new dimension: the cloud becomes a means by which enterprises can share business logic.

Cloud ecosystems foster standards-based business exchanges between participating enterprises. This encourages optimization of the products and services available in the ecosystem's marketplace.

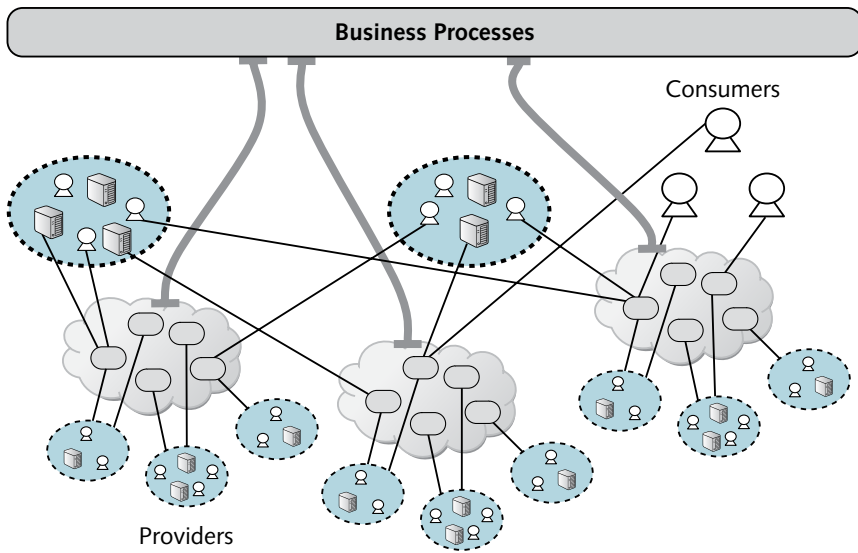


Figure 1.10: Cloud ecosystem

Cloud computing enables providers and consumers of products and services to interact much more easily, because they are freed from the drudgery of providing the necessary IT support. This encourages the creation and growth of ecosystems in which companies co-operate effectively to meet the needs of society. The benefits of cloud computing extend not just to individual businesses, but to business as a whole.

