THE OPEN GROUP SERIES

Cloud Computing for Business

The Open Group Guide

Chris Harding and members of
The Open Group Cloud Work Group

Van Haren Publishing

THE Open GROUP

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Other publications by Van Haren Publishing

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

- 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 material etc. in the **VHP Knowledge Base**: www.vanharen.net for more details.

VHP is also publisher on behalf of leading organizations and companies:
ASLBiSL Foundation, CA, Centre Henri Tudor, Gaming Works, Getronics, IACCM, IAOP, IPMA-NL, ITSqc, NAF, Ngi, PMI-NL, PON, Quint, The Open Group, The Sox Institute

Topics are (per domain):

IT (Service) Management / IT Governance

ABC of ICT
ASL
BiSL
CATS
CMMI
CoBIT
ISO 17799
ISO 27001
ISO 27002
ISO/IEC 20000
ISPL
IT Service CMM
ITIL® V3
ITSM
MOF
MSF
SABSA

Architecture (Enterprise and IT)

Archimate®
GEA®
SOA
TOGAF®

Business Management

CMMI
Contract Management
EFQM
eSCM
ISA-95
ISO 9000
ISO 9001:2000
OPBOK
Outsourcing
SAP
SixSigma
SOX
SqEME®

Project/Programme/ Risk Management

A4-Projectmanagement
ICB / NCB
MINCE®
M_o_R®
MSP™
P3O
PMBOK® Guide
PRINCE2®

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Cloud Computing for Business

The Open Group Guide

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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
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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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 ***versus*** 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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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 *versus* 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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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 *versus* reward 49
Figure 3.2: Impact on products and services 50
Figure 3.3: Business operating scope *versus* complexity 53
Figure 3.4: Collaboration *versus* information restriction 54
Figure 3.5: Deciding on your cloud solution 58
Figure 3.6: Konsort-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: Konsort-Prinz capacity utilization 104
Figure 4.9: Konsort-Prinz peak day utilization 104
Figure 4.10: Konsort-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: Konsort-Prinz initial risk assessment 126
Figure 5.6: Konsort-Prinz risk mitigation 127

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Figure 5.7: Konsort-Prinz supplier cost models 128
Figure 5.8: ViWi six-month risk assessment 129
Figure 6.1: The capacity *versus* 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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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."

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

"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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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).

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

- Karin Middeljans, Netherlands Ministry of Economic Affairs, Agriculture and Innovation
- Vinod Ralh, IBM
- Chiel Steinvoort, 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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

- 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 Riousset, IBM
- Michael P. Rogers, Capgemini
- Rajib Roy, HP
- Tarun Roy, Capgemini
- Nayan Ruparelia, HP
- Ilyas Sener, Capgemini

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

- 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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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/dorofeetutorialndia09_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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

[EUCALYPTUS] Eucalyptus: An Open Source Cloud Computing Infrastructure, Daniel Nurmi, Rich Wolski, Chris Grzegorczyk, 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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

[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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

[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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

- 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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

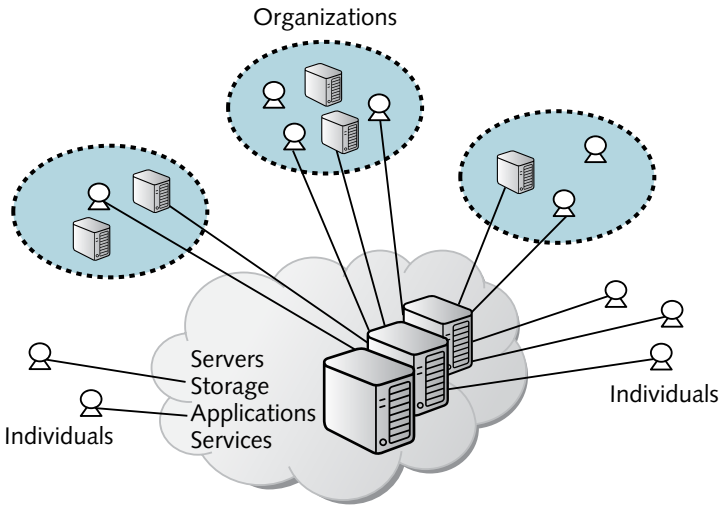

Figure 1.1: Basic cloud computing model

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).

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

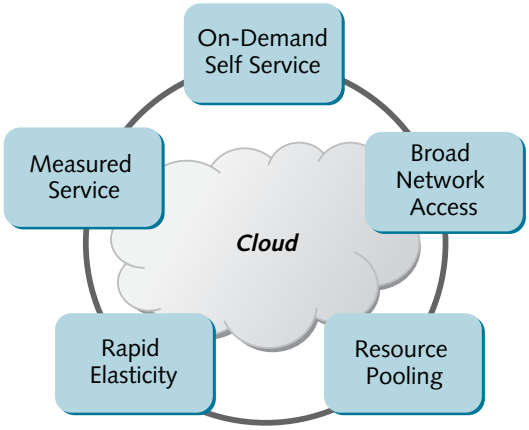

Figure 1.2: The essential characteristics of cloud computing

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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."

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

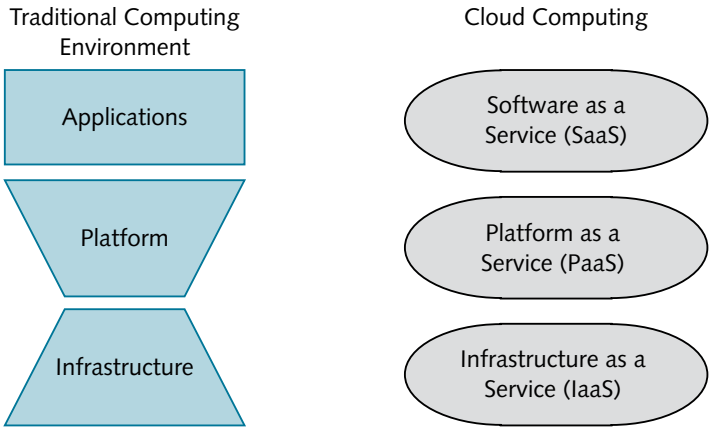

Figure 1.3: Cloud service models

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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."

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

- 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."

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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).

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

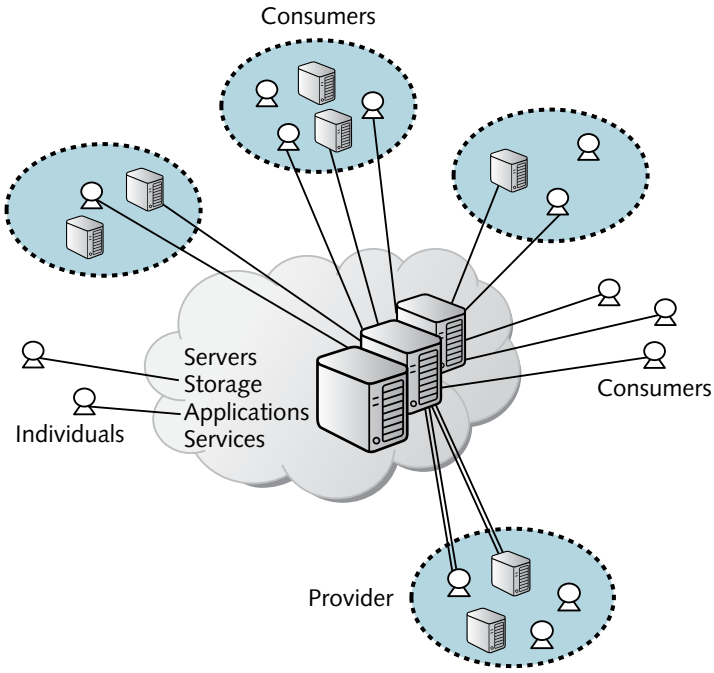

Figure 1.4: Providing cloud services

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

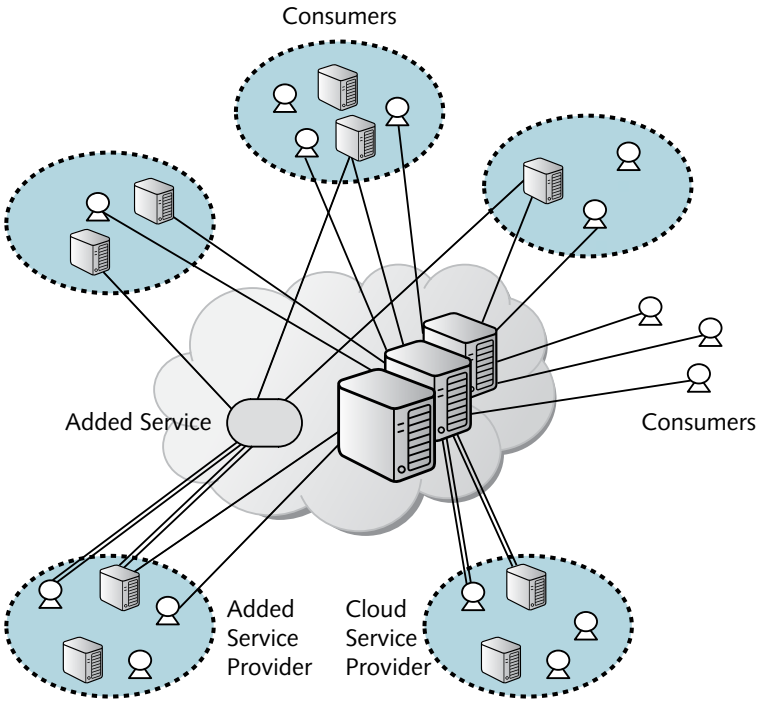

Figure 1.5: Providing added services

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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

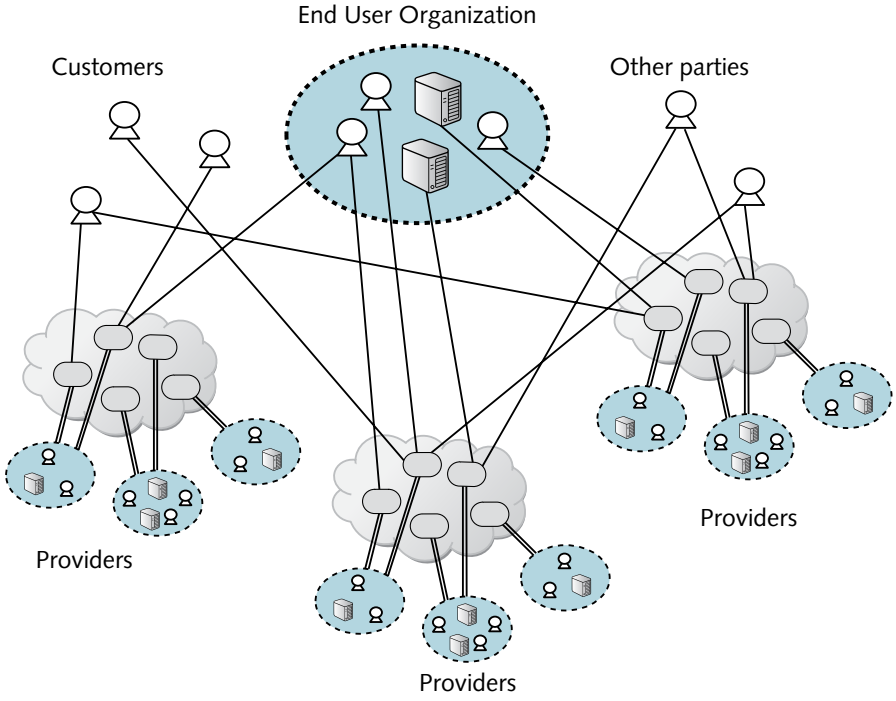

Figure 1.6: Using cloud services

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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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 polices 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

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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 ceate 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).

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

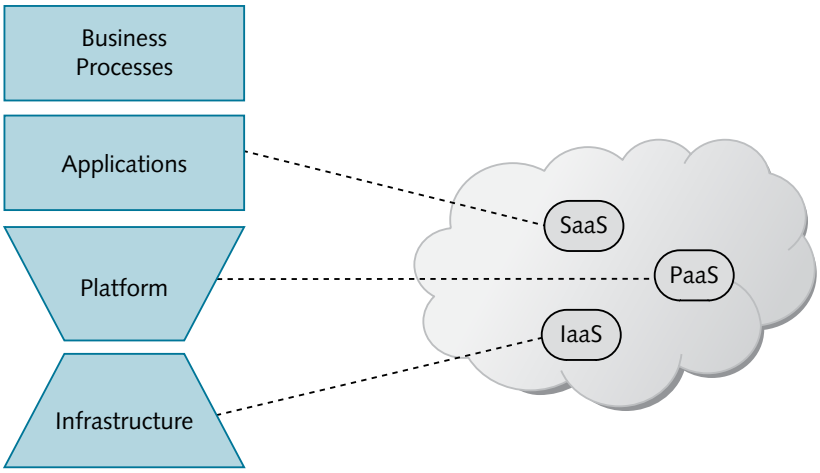

Figure 1.7: Use of cloud services to support business processes

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.

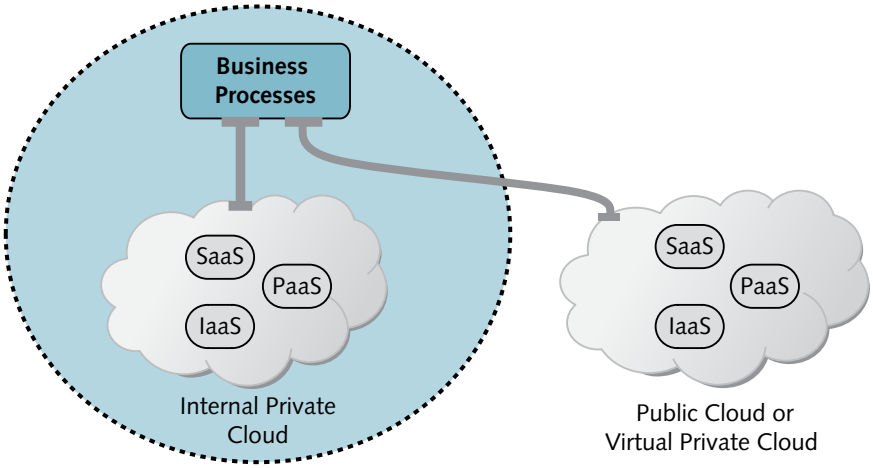

Figure 1.8: Internal and external cloud services

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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.

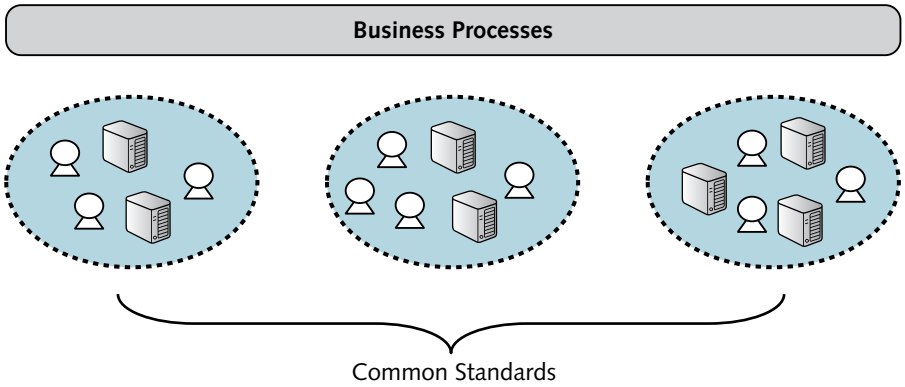

Figure 1.9: Business ecosystem

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

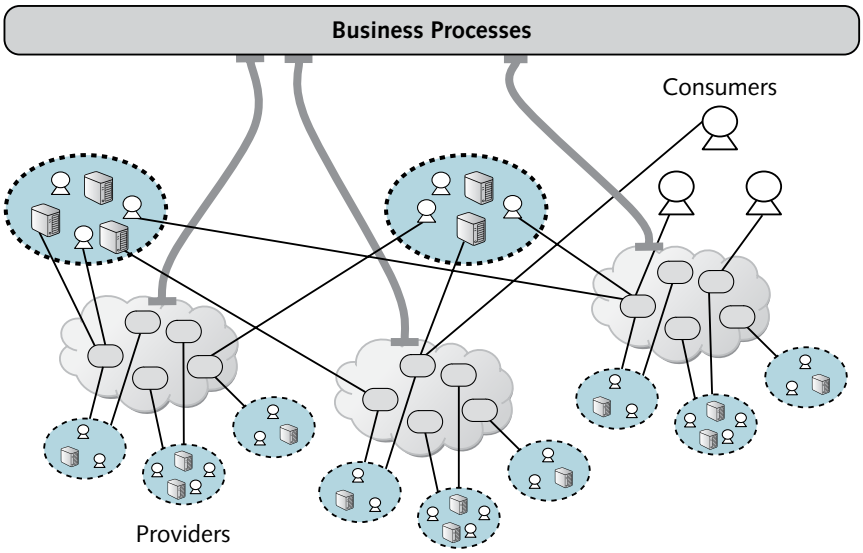

Figure 1.10: Cloud ecosystem

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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Chapter 2

Why cloud?

There are many reasons why an enterprise might want to adopt cloud computing. The Open Group Cloud Business Use-Cases project analyzed 24 business use-cases. They are listed in Cloud Computing in Use (Appendix A). They form a good representative sample, illustrating the variety of reasons why enterprises are turning to the cloud.

The main business elements emerging from the use-cases are shown in the figure below (Figure 2.1).

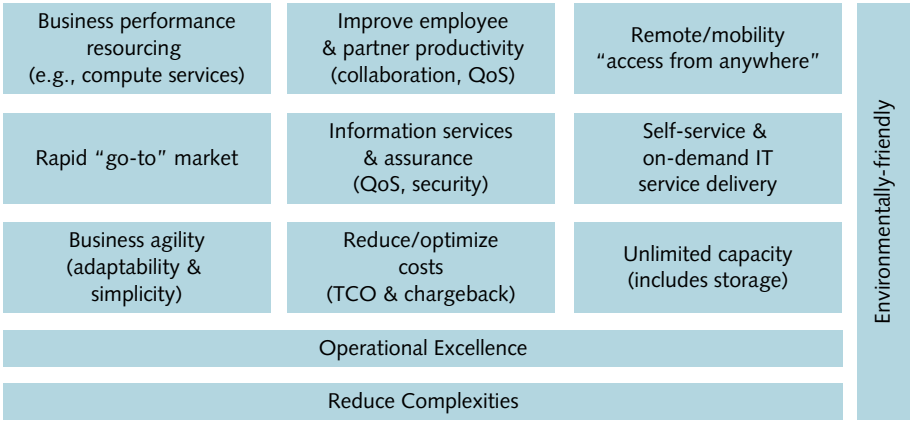

Figure 2.1: Use-case business elements

The most significant ones are identified as business performance resourcing, rapid go-to-market, business agility, and cost reduction.

Often, there is no single reason why an organization chooses to adopt cloud computing. The decision depends on a complex combination of reasons, rather than being based on a single factor.

The reasons do not apply universally. Cloud computing usually does provide greater agility, for example, but it does not always reduce costs or mean a better Quality of Service (QoS). The Cloud Buyers Decision Tree (see Section

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

3.3) will help you to assess whether cloud is appropriate in your situation and, if so, what kind of cloud service to choose. Building ROI from Cloud Computing (Chapter 6) discusses the factors of cost, agility, and quality in more depth.

There are a number of concerns that must be addressed when architecting a cloud solution. These concerns can often be addressed, but in some cases they may mean that cloud is not appropriate.

Frequently expressed concerns include:

- Security, privacy, and compliance, including data ownership and availability in shared environments, regulation and legal issues, corporate policies, and identity management for access control
- Lack of functionality or inflexibility of SaaS applications
- Dependency on an Internet connection
- Vendor lock-in as a result of lack of standards, and portability issues due to immaturity of cloud products
- Vendor management, which requires a different approach, in which SLAs are critical
- Change management and testing, which can be a challenge, especially in shared environments
- Integration with on-premises systems, which may be difficult, or even impossible
- Lack of transparency of interfaces between SaaS vendors, particularly with regard to managing the interfaces
- Lack of experience with cloud financial models and licensing

Use of cloud computing can change a company's risk posture and profile. Use of public cloud, for example, can mean that you avoid the risk that a large investment in IT resources will not pay off but, on the other hand, it can introduce security risks, due to sharing resources with unknown parties, and risk of non-compliance with regulation, due to lack of control over how and where data is stored. The change can be positive or negative, depending on your circumstances. See Understanding Cloud Risk (Chapter 5) for a discussion of how to assess cloud-related risks.

This chapter describes the main reasons for using cloud computing, and discusses the basic situations in which they may or may not apply. It is based

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

on the use-cases developed by The Open Group Cloud Business Use-Cases project and includes some considerations identified in The Open Group Cloud Report on the Computing Business Scenario Workshop [BUSSCEN]. It will help you start to gain the understanding that you need to obtain the benefits that cloud computing can bring.

The reasons can be classified under five topics: agility, productivity, QoS, cost, and the ability to take advantage of new business opportunities.

2.1 Agility

The first main reason for using cloud computing is agility. Cloud computing delivers improved agility because it has on-demand self-service and rapid elasticity. IT resources can be acquired and deployed more quickly and, once deployed, they can be increased or decreased as needed to meet demand. This means that enterprises can innovate, introduce new products and services, enter new markets, and adapt to changing circumstances.

This applies to all cloud deployment models. For public cloud, there is an additional consideration. Because the cloud resources are operated by an external organization, there is less need to train people within the enterprise. Skills acquisition often takes longer than any other operation associated with deploying a new resource. The reduced training needs for public cloud are a major contributor to cloud-based agility.

2.1.1 Changing business processes

Business agility requires the ability to create new business processes and change existing ones. This often means adding to or changing the supporting IT resources. The difficulty of doing this, and the time that it takes, can be a major barrier to innovation. Cloud computing can remove this barrier by enabling the enterprise to add to or change its IT resources quickly and easily.

In the cloud business use-cases investigated by The Open Group, Case 19: Rapid Deployment is an example of an organization that is looking to cloud computing to support new services, and Case 22: Innovation – New Ventures is a company that is looking to cloud computing to help it streamline its processes and be more responsive to its customers. In Case 7: Business Transaction Assurance and Continuity, frequent mergers and acquisitions

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

mean drastic changes to business processes, which cloud computing can make less painful.

Both public and private cloud can deliver this benefit. Any enterprise can use public cloud to obtain new IT capacity quickly. In an enterprise that has an established private cloud, the divisions and departments can use it in the same way.

2.1.2 Development and testing

As well as resources to support the new business processes themselves, IT resources are often needed to support development and testing of new software. Procurement of these resources can introduce significant delay. In addition, the need to pay for expensive resources for development of a new product or service can add significantly to the business risk.

Public or private cloud can eliminate the delay by enabling development resources to be available on-demand. Public cloud can reduce the cost of development and eliminate the additional business risk because it replaces an up-front investment in development resources by usage-based payment.

These reasons – faster development and reduced up-front investment – are among those most frequently given by companies considering cloud computing. They appear in Case 8: Business Market Launch, Case 20: Rapid Business Process Extension, and Case 22: Innovation – New Ventures. In Case 6: Development and Test Services, Case 15: Business Growth, Development Test, and Validation Support, and Case 18: Test Assurance Services, the organizations concerned are specifically looking for development and test platforms, and see cloud computing as a good way of providing them.

2.1.3 Resource scaling

Deployed services may experience increases and decreases in usage volumes that require corresponding increases and decreases in supporting IT resources to meet user demands in a cost-effective way. These increases and decreases can be handled easily by using cloud-based resources.

The Konsort-Prinz fictional example in this Guide illustrates many of the considerations involved. Examples of real situations where this is important

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

are provided by Case 3: Rapid Business Capacity and Scale, Case 10: Research, and Case 12: Business Utility/Agility.

In some cases, resource provisioning is possible using programmatic interfaces, so that resources can be scaled automatically under program control without human intervention. This can cut the time to react to change from hours or days to minutes.

Resource scaling, whether manual or automatic, enables service levels to be maintained and reduces cost. Automatic resource scaling matches resources to demand most efficiently, enabling the lowest cost of operation consistent with adequate service.

2.1.4 Reduced need for training

As with other forms of outsourcing, public IaaS, PaaS, and SaaS enable an organization to acquire operation and support capabilities for IT infrastructure, platforms, and applications, without its staff having to acquire the corresponding specialist operation and support skills. This can dramatically cut the time taken to deploy these capabilities. In some cases, it can enable the organization to acquire capabilities that it would otherwise not consider.

The ability to acquire capabilities without the need for training is an important consideration in Case 2: Modernization, Case 9: Secure Business Assurance Services, Case 11: Productivity, Case 14: End-User Cost Optimization, and Case 15: Business Growth, Development, Test, and Validation Support.

Note that use of public or private cloud may require new service management skills, and use of private cloud may require new service provision skills. Some training in these skills may be needed. Use of cloud computing does not eliminate the need for training entirely. It can remove the need for training in specific products and services, but may introduce a need for generic training that applies to a range of cloud-based services.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

2.2 Productivity

Increased productivity is the second main reason for using cloud computing. It provides a more productive environment for collaborative working, and improves productivity by enabling participants in a business ecosystem to share processing logic.

2.2.1 Collaborative working

Use of cloud-based tools for email, instant messaging, voice communication, information sharing and development, event scheduling, and conferencing is becoming an increasingly common feature of business life. This is so much the case that people often do not think about it as being 'cloud computing'. Nevertheless, these services generally exhibit the five essential cloud computing characteristics, and are part of the growing trend to do business 'on the cloud'.

A particularly important characteristic for these services is Broad Network Access (BNA). It is widely used on mobile phones and browser-less devices as well as on PCs with web browsers.

People from the so-called 'millennium generation' are starting in business life today. They are accustomed to using cloud services for social interaction and sharing personal information. They expect to have business collaboration tools that are equally effective and easy to use.

Increased productivity of collaborative working is a reason for using cloud computing in Case 5: People Productivity, Workplace Collaboration, Case 17: Partner Collaboration, Case 19: Rapid Deployment, and Case 21: Brand Unification.

2.2.2 Shared logic

The ability of the cloud to provide shared logic in a business ecosystem was discussed under Cloud Ecosystems (in Section 1.8).

The desire for this is shown in Case 12: Business Utility/Agility, where the ecosystem is a government agency, and the requirement is for common utility services to support inter-departmental activities.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

As collaborative working between organizations becomes increasingly important, this capability of cloud computing will be expressed as a requirement increasingly often.

2.3 Quality

The third main reason for using cloud computing is that it can mean better quality IT. There are several reasons for this. How far they apply varies from one enterprise to another: an enterprise with an inefficient IT department is likely to find cloud a more attractive proposition than one whose IT department provides a superb service, for example.

These reasons include better usage information, better manageability, better quality of IT provision, better business continuity, and better carbon footprint.

Some of them apply particularly to public cloud rather than private cloud, and may equally apply to other forms of outsourcing. If you are considering public cloud in order to improve the quality of your IT, you should also look at other forms of outsourcing.

2.3.1 Better usage information

Measured Service is one of the essential characteristics of cloud computing. In a public cloud, it is an essential basis for charging. In a private cloud, it may be used for internal charging, and should in any case be used for resource planning.

There is no fundamental reason why service use cannot be measured without using cloud computing, but often it is not measured sufficiently. The virtualization technology that underpins many cloud solutions provides convenient measurement facilities that other technologies may lack. Good usage information is a given for cloud, but not for other kinds of solution.

The value of better usage information for IT management purposes is shown by Case 10: Research and Case 13: IT Cost Optimization. In Case 1: Business Activity Reporting, good usage information delivers a direct business benefit.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Good usage information provides enterprise stakeholders with an understanding of how the enterprise IT is operating, and enables effective forward planning, equitable sharing of resources, and more efficient resource use.

2.3.2 Better manageability

Better manageability is a reason for adopting cloud in Case 4: Operational Efficiency and Case 10: Research.

The ability to provision and configure resources through a web portal – or even automatically – means that consumers of cloud services have substantially better manageability than they would expect to have with a non-cloud system.

For providers, the fact that resource provisioning and configuring is done automatically in response to user requests means that a heavy management burden is removed.

2.3.3 Better quality of IT provision

An enterprise can very easily be 'locked in' to its internal IT department, and unable to drive up quality by using competition as a lever, as it would for external suppliers. Where an in-house department does not keep abreast of the latest developments, adopts a 'not invented here' attitude, or becomes ineffective or inefficient in other ways, it can be hard to remedy the situation. Such problems often first become visible in the help desk, but this may be only symptomatic of a deeper malaise.

Ineffective or inefficient IT can severely compromise an enterprise's competitiveness. IT is not a core competence of most enterprises, but their dependence on it forces them to give it significant attention and management resource, which would be better devoted to their core activities.

By using public cloud services, you can reduce your dependence on internal IT, reduce the need to devote management attention and resource to your IT infrastructure, and use the competitive dynamics of the cloud marketplace to obtain an IT infrastructure that is effective and efficient.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

With public or private cloud, the ability of users to obtain and configure resources by self-service can be a significant quality improvement, regardless of the capability of the IT department.

2.3.4 Better business continuity

Businesses wish to continue their operations as normally as possible in the event of disaster, such as a fire or flood destroying a data center, but the cost of providing duplicate systems and recovery mechanisms in-house can be very high. This is a significant pain-point, especially for small businesses, where the cost is proportionately higher than it is for large ones.

This means that you must either incur significant costs that decrease your competitiveness, or risk failing when disaster strikes.

Cloud can solve this problem. For large cloud vendors, economies of scale make the cost and effort of providing duplicate systems and recovery mechanisms a relatively small overhead, so that they can deliver disaster recovery capabilities to businesses of all sizes at low cost. Alternatively, businesses can implement disaster-recovery capabilities themselves by using cloud services whose underlying resources are geographically separated.

Cloud computing is considered as a means to ensuring business continuity in Case 7: Business Transaction Assurance and Continuity.

2.3.5 Better carbon footprint

In the current social and political climate, many companies want to do all that they can to increase their 'green' (i.e., environmentally-friendly) credentials and reduce the carbon emissions of their operations. But they are expecting to see rapid growth in their use of IT, and corresponding growth in emissions from increased power consumption.

Through server consolidation, optimal resource utilization and use of thin clients (see below), cloud computing can make more efficient use of resources. This means a corresponding reduction in carbon footprint. With public cloud, there is the added advantage that the resource used is off-premise.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Case 23: Green Footprint in Cloud Computing in Use (Appendix A) is an example of a situation where companies are looking to cloud computing to help them reduce their carbon footprints.

2.4 Cost

Given the advantages of agility, productivity, and quality that cloud computing can have, you might expect it to be generally more expensive. This is, however, not always the case, and reduced cost is one of the main reasons why companies are turning to the cloud.

Cloud computing can help you achieve cost reduction through server consolidation, thin clients, or community cost sharing.

Another reason related to cost is the ability, with public cloud, to replace capital expenditure (CAPEX) by operational expenditure (OPEX) in some cases.

2.4.1 Server consolidation

The traditional IT model leads to the provisioning of servers to match peaks of demand. There is also a tendency for servers to be dedicated to particular departments or projects, with no attempt to balance surplus in one area against shortage in another. The result is often a massive overcapacity, with very poor utilization of resources.

The Federal Cloud Computing Strategy published in February 2011 [FEDSTRAT] describes how US federal government IT suffers from inefficient use of resources. It estimates that, by using the cloud computing model for IT services, the US federal government will be able to reduce its data center infrastructure expenditure by 30%.

Reduced cost through resource pooling has long been one of the arguments for outsourcing. But moving the responsibility to an external supplier is not the fundamental point. A large enterprise might just as well achieve cost savings by optimizing server utilization across departments – yet few do.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Until recently, this kind of optimization would largely have been done manually. This is difficult and time-consuming, and the delay between the change in resource use and the corresponding re-allocation of resources makes manual optimization far from perfect. Virtualization technology, with software-driven resource allocation, makes the process automatic. This technology is commonly used by cloud suppliers. It means that cloud computing delivers fast and effective resource optimization, moving processor, memory, storage, and network capacity between consumers almost instantaneously.

In the case of private cloud, the enterprise installs the virtualization and resource allocation technology as part of the cloud solution, and then obtains cost savings through improved resource utilization. With public cloud, the enterprise obtains cost savings by paying less to the cloud provider than it would have paid to run its own systems inefficiently. In this case, the cloud provider is likely to keep a proportion of the savings, as a fair return for its services. Private cloud may therefore be the cheaper option overall, particularly for a large enterprise.

You should consider whether the overall traffic load fluctuates widely, or is relatively constant. If there are large variations in your resource needs then public cloud can deliver cost savings by pooling your traffic with that of other users. If, however, your needs are pretty constant then you will be paying a premium for elasticity that you do not use when you buy public cloud services. If your overall load is constant but composed of a number of variable project or department loads, then a private cloud can reduce your costs. If you can separate your load into constant and variable parts, hybrid cloud could be your best choice. If all of your component loads are constant, or vary predictably, and you can optimize manually, cloud will not produce a saving for you.

The fictional Sam Pan Engineering example in this Guide illustrates how cloud computing can reduce cost through server consolidation. Case 10: Research and Case 13: IT Cost Optimization are real examples where efficient use of resources is important.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

2.4.2 Thin clients

Another option for cost reduction is to move processing from clients to cloud-based servers. The 'thin client' is not a new idea, and can be used with traditional solutions. It is, however, often classed as a form of cloud computing because the processing is moved from the client 'to the cloud'.

Case 14: End-User Cost Optimization and Case 19: Rapid Deployment are examples of situations where this kind of solution would be considered.

Significant cost reductions can be obtained by replacing expensive client devices – such as high-specification PCs with high processing power and large amounts of memory – by cheaper, less powerful client devices that just provide a user interface to server applications that do most of the processing. This is, however, likely to mean a substantial increase in traffic between client and server across the network, and will not be viable unless the network is reliable and fast.

2.4.3 Community cost sharing

Community cloud provides a way for a community of enterprises to share costs of common resources. If viable payment and governance models can be established – and these may not be easy things to achieve – then community cloud can be a very cost-effective way for those enterprises to obtain IT resources.

Case 24: Virtual Learning provides an example of community cost sharing.

Community cloud is perhaps the deployment model that is least talked about, at least in a commercial environment. It may nevertheless have a significant role to play, just as community software development has come to play a significant role through the open source movement.

2.4.4 Replacing CAPEX with OPEX

A company with a high cost of capital, or which would benefit from reducing its tax burden, is a candidate for shifting CAPEX to OPEX, but other aspects of the business context, such as availability of appropriate solutions, or security constraints on using shared services, may contradict that candidacy. If a move from CAPEX to OPEX is desired, then the company should be

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

considering and evaluating outsourcing solutions, including public cloud, hybrid cloud, and outsourced private cloud.

Cash flow can be an important indicator if CAPEX to OPEX is the focus. Pay-as-you-go can be seen as easier on cash flow than pay-upfront. But both cash flow considerations may not necessarily exist in the same scenario. For example, a business may want to improve cash flow by moving to a direct usage model but still retain investment in CAPEX for differentiated private business processes.

Using an OPEX model can potentially release capital that would otherwise be used for initial investment and ownership of IT assets. Alternatively, you could invest in a cloud computing platform and make changes to the payment and funding of the service as it is amortized over a wider shared service model for economies of scale.

The cost of capital from sources of equity and cost of debt point of view can change for private and public industries that have stock market and government sources of funding.

If the overall goal is to maximize the use of capital by best use of the debt and equity funds, in cloud computing the use of OPEX moves the funding towards optimizing capital investment leverage and risk management of those sources of funds.

2.5 New business opportunities

Cloud computing can give an enterprise new business opportunities as a provider of cloud services or added services.

2.5.1 Cloud service provision

A company that excels in the quality of its IT can become a public IaaS or PaaS provider.

A particular case of this might be where a company implements a private cloud, has spare capacity, and sells that capacity as public cloud.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Processing power, data storage, and computer operating system capabilities form the basis of most IaaS and PaaS services today. There may be opportunities to sell infrastructure or platform services based on other kinds of resource. For example, in Case 16: Service Delivery, a telecommunications service provider considers how it could provide 'network as a service' and 'transport as a service' to online consumer service companies.

Software applications providers may well find that they can increase the markets for their products by providing them in the form of SaaS.

2.5.2 Added service provision

The provision of added services using cloud services is described under Providing Added Services (in Section 1.7). By implementing your services on the cloud, and making them accessible via the Internet, you can reach a vast, global market.

As cloud computing grows in scope and popularity, it will increasingly become the favored way for enterprises to deliver their services. Competition will be fierce, and only the best will succeed. For those that do succeed the rewards will be high.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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 – such as cloud computing – as the way to solve your problem, or seize your opportunity. This is your architecture vision.

An IT architecture vision shows how a business need can be met, or opportunity realized, by a particular combination of technology capabilities. It provides the sponsor with a key tool to sell the benefits of a proposed project to stakeholders and decision-makers within the enterprise, or to external investors.

Case 5: People Productivity, Workplace Collaboration in Cloud Computing in Use (Appendix A) sets out a vision: "We will improve communication within and outside the organization by using cloud-based collaboration tools." For Case 7: Business Transaction Assurance and Continuity, the vision is: "We will ensure that our business can continue in the face of disruption or disaster by using back-up resources on the cloud." Each of the business use-cases in that appendix encapsulates an architecture vision that has cloud computing as its technological basis.

This chapter describes how to establish such a vision. It discusses how to assess your business context, and determine whether cloud computing will help you reach your goals. This discussion is illustrated by three examples, which are introduced here and will be used throughout the rest of the book.

This Guide is concerned with IT solutions that are solutions to business problems, and that use cloud services. A complex solution of this kind should not be built *ad hoc*; it should have a properly thought-through architecture. An architecture vision is the starting point for the development of such a solution architecture.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

For the enterprise as a whole, an overall enterprise architecture will define roadmaps for the development and deployment of solutions. While cloud computing may not be the basis of your enterprise architecture, you should consider it as a significant component, and develop a cloud roadmap to ensure that cloud computing is incorporated in your solutions in a sensible and co-ordinated fashion.

This Guide does not address the further process of developing enterprise or solution architectures. Refer to The Open Group Architecture Framework [TOGAF] for more information on architecture development.

3.1 Understanding the business context

When looking at the technological possibilities, you must fully understand the business context. This section discusses the aspects of your business context that affect the decision to use cloud computing.

3.1.1 Basic situation

Consider the size of your enterprise, the industry sectors in which it operates, its geographical distribution and locations, and legislation and other external regulation that applies to your business.

3.1.2 Business goals

What are the business goals you seek to address?

As described in the previous chapter (Chapter 2), you might be looking to improve your agility, improve your productivity, improve the quality of your IT, improve your cost model, or take advantage of new business opportunities.

You should analyze and describe your business goals as the starting point of your understanding of your business context.

3.1.3 Risk-reward balance

What is the acceptable risk-reward balance?

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

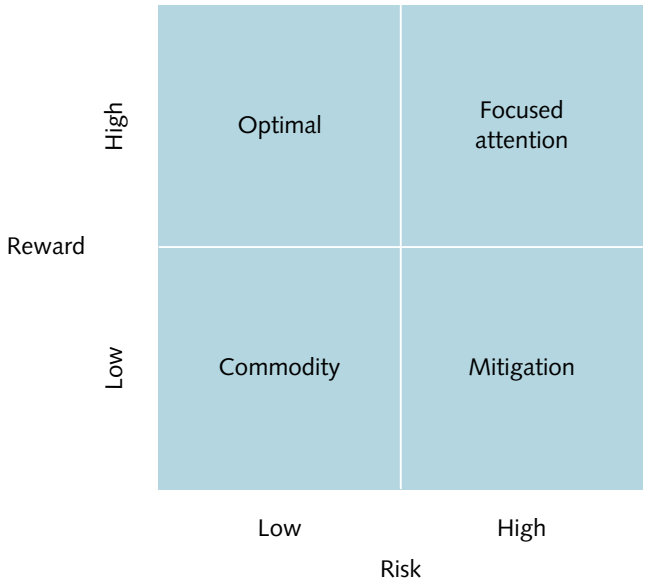

Figure 3.1: Risk *versus* reward

For low risk, consider whether low-risk activities can be moved to an on-demand service environment, and whether activities can be identified to explore and exploit on-demand business opportunities with mitigated risk management through partnerships and risk sharing with cloud services.

For high risk, consider:

- What business process operations carry high corporate risk in operational and legal failure
- Whether high corporate risk can be mitigated further by enhanced risk sharing and disaster recovery and business continuity practices supported by on-demand back-up and recovery services
- Whether there are specific corporate risks that can be linked to specific IT software applications, infrastructure, and services
- Whether corporate benefits should be traded-off against corporate risks

For low reward, consider whether business activities with low corporate return can be commoditized for low-cost competitive advantages.

For high reward, consider whether there are opportunities to improve market share, revenue, profit, or cost management through on-demand delivery of cloud services.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

3.1.4 Impact on products and services

Does the business seek to explore new products and services in new markets (transformative and disruptive), or to exploit existing markets with current products and services (utility and commodity)? This positioning is a key aspect of business strategy and delivery.

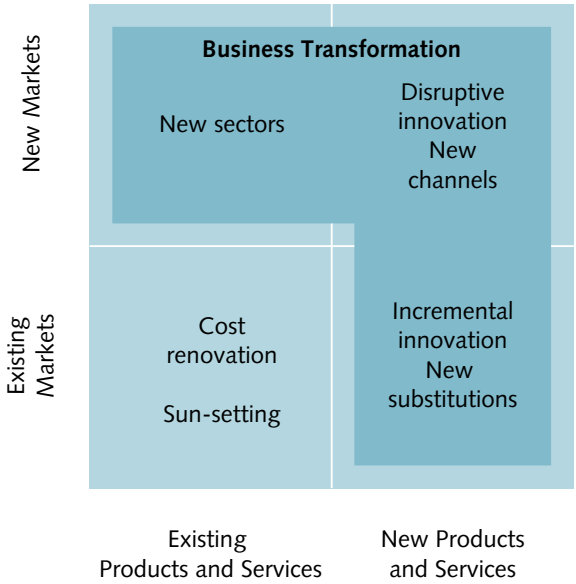

Figure 3.2: Impact on products and services

Existing products and services can be offered through on-demand self-service, augmented and enhanced through on-demand features and options, or commoditized for competitive low-cost advantages.

New products and services can be sourced and delivered on-demand. They can be disruptive to existing products and services through superior cost and functional performance, or offered as commodities for competitive low-cost advantages.

New and existing markets can be expanded through rapid scaling and expansion of products and services offered, or entered by offering commodity products and services for competitive low-cost advantages.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

3.1.5 Business processes

Which business processes are under consideration?

You may be considering:

- Your complete business process portfolio
- No specific business processes
- One or more specific processes
- The processes in one or more specific business areas

You should identify the specific processes or specific process areas. Example business areas and activities carried out by business processes are shown in the table below (Table 3.1).

Business Area	Business Activity
Human Capital Management (HCM)	Skills Management Compensation Planning Benefits Administration
Financial Management (FM)	Accounting: General Ledger, Accounts Payable, and Accounts Receivable Payroll Taxes Order to Cash Real Estate Management Financial and Compliance Reporting, Sarbanes-Oxley [SOX], and BASEL II Business Performance Management Risk Management
Customer Relationship Management (CRM)	Sales Customer Experience Management (CEM) Call Center Management Sales Force Automation (SFA) Business Intelligence (BI) Business Analytics Campaign Management Sales Analytics
Supply Chain Management (SCM)	Procurement Inventory Management Import Compliance Supplier Relationship Management (SRM) Logistics
Manufacturing	Product Lifecycle Management (PLM) Workforce Management Resource and Capacity Management

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Business Area	Business Activity
Information Technology (IT)	IT Architecture Software Development Data Center Operation
Corporate	Research & Development (R&D) Strategy and Portfolio Management Marketing Communications Legal

Table 3.1: Example business processes

Consider whether these business processes are differentiating, part of your company's competitive advantage. Differentiating business processes are the only processes for which a case for in-house, custom, dedicated enablement independent of the rest of the business portfolio can readily be made.

For non-differentiating, 'business-as-usual' processes, such as benefits, inventory, accounting, and intellectual property management, you should seek to minimize Total Cost of Ownership (TCO), while maintaining functionality and Quality of Service (QoS). For non-differentiating processes concerned with external compliance, such as SOX II compliance, ethical procurement, and carbon footprint management, you should seek to achieve timely delivery of functionality, to meet QoS requirements, and to minimize TCO.

If your processes are part of your competitive advantage, consider whether the differentiation is based on IT capabilities. An example is information retrieval supported by an Internet search engine. In such a case, you should seek to align with the corporate mission and strategic objectives, and erect barriers such as intellectual property protection to prevent imitation. If the IT is not a differentiating factor, as for example with online sales and store returns, you should seek to meet functional and QoS requirements, and adapt quickly to changes in those requirements.

Consider how far the IT resources used by the business processes are dedicated and customized. IT resources include software applications, IT infrastructure, and development and support personnel. They may be fully dedicated to the processes under consideration, partially dedicated, or fully shared. They may be customized, tailored, or standard.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Consider also the degree of dedication and customization of the non-IT resources used by the processes. These include labor, fixed assets, and process definitions.

3.1.6 Scope and complexity

What is the scope of operation relative to the process complexity?

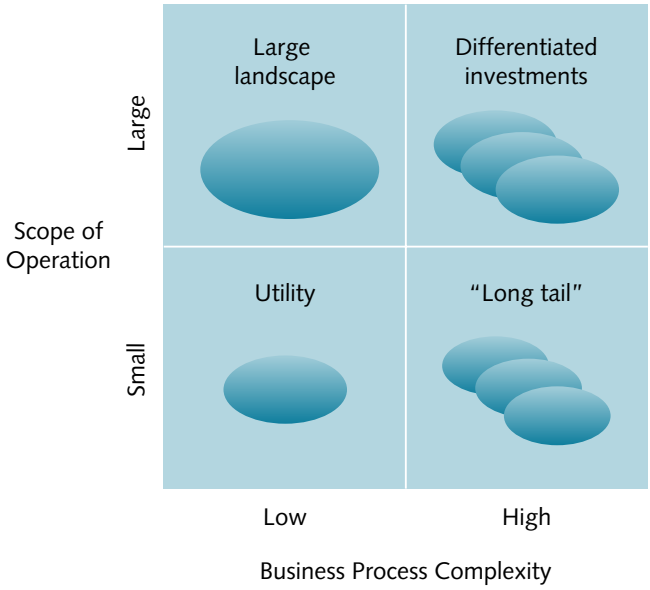

Figure 3.3: Business operating scope *versus* complexity

For a small scope of operation, consider what niche business process operations could be moved to an on-demand hosting provider, whether specific IT operations can be moved to an on-demand hosting provider, and whether specific IT operations can be commoditized for competitive low-cost advantages.

For a large scope of operation, consider what scope of existing business operation can be moved to an on-demand hosting provider, and whether a large scope of existing IT landscape can be moved to an on-demand model.

For low business process complexity, consider whether low-complexity business processes can be commoditized for low-cost competitive advantages.

For high business process complexity, consider what specific business processes need to be retained under control of the business for competitive

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

advantage, and whether highly complex business processes can be improved through complexity reduction.

3.1.7 Collaboration *versus* information restriction

What is the balance between collaboration and the need to restrict information?

You should understand the degree of collaboration and types of data that can be moved into a cloud environment, and assess the different cloud deployment and hosting models that meet their sovereign, corporate, and personal security and collaboration needs.

The degree of collaboration can range from individual working to integrated community activity. The information involved can be private to the individual or community, or can be public.

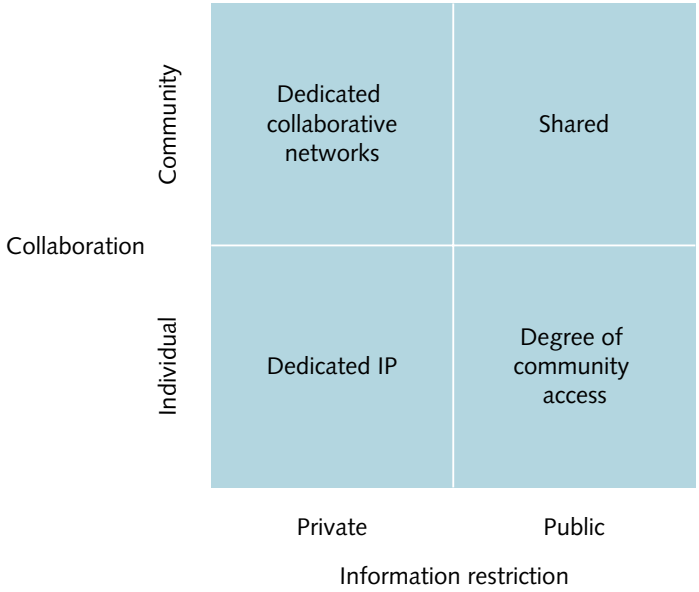

Figure 3.4: Collaboration *versus* information restriction

For individual operation, consider whether desktop and mobile working can be improved through on-demand personal productivity tools (such as email, bulletin boards, and office applications), and whether the personal information and assets created by individuals are the private intellectual property of the corporation, or can be shared.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

For community collaboration, consider whether business participants and stakeholders can collaborate in a shared platform and business service environment, and whether the information and assets created by collaboration are the private intellectual property of the corporation or collaborating group, or can be shared.

For private information, consider whether current corporate and personal data rules identify and define who owns what intellectual property, and whether corporate and personal data can be partitioned and isolated effectively for secure storage and access control.

For public information, consider whether:

- Current corporate and personal data rules prohibit the storage and access to personal and corporate information in public locations
- Legislation (such as EU Data Protection legislation or the US Patriot Act) affects personal and private data held in public locations
- Information held publicly can be monitored and managed to e-Discovery standards for legal purposes
- Personal and corporate information can be selected, partitioned, and isolated effectively for storage and use in public locations

3.2 Three example cloud projects

In the rest of this Guide, we will use three fictional companies as examples to illustrate the ideas and bring the explanations to life. The descriptions of them are simpler than you might expect for a real-life company: this makes it easier to bring out the key points without being confused by extraneous detail. They are chosen to show different aspects of cloud computing. They are loosely based on the use-cases described in Cloud Computing in Use (Appendix A), but do not correspond to particular cases. The authors apologize if any real company is offended by any accidental resemblance.

3.2.1 Konsort-Prinz

This is a mid-sized European company that sells seasonal novelties by mail order. It started in the 1950s with Christmas decorations, and has since developed products for other festivals of various kinds, but the bulk of its business is still in December, with a few regular high spots during the rest of the year, and occasional, unpredictable 'one-offs' (such as royal weddings).

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Everything it sells is bought from small manufacturers; it produces nothing itself. It uses in-house IT for its website and for order-processing, supplier management, and stock management. This is substantially under-utilized for much of the time, and Konsort-Prinz is considering using cloud computing to reduce its capital costs and obtain a better return on its investment.

This company:

- Is looking to scale its IT resources appropriately
- Is looking to obtain a good reward, but at a low risk
- Will just offer its existing products and services, and in its current markets
- Is considering cloud for its website and for order processing, supplier management, and stock management primarily (but could also use it for other business processes)
- Is looking at a large scope of operation and process complexity
- Has some collaboration with suppliers, but will keep information private within that community

3.2.2 Sam Pan Engineering

This is a large conglomerate headquartered in Hong Kong and spread over South-East Asia and Australasia. It has a range of operations, from making small pleasure-yachts to installing water-purification equipment and air conditioning. It has factories and offices in many cities in several different countries. At present, most of these have some kind of IT capability, but this is very inefficient, and the company has decided to establish area data centers that will offer IT capabilities to the scattered locations. It believes that it could implement this as a private cloud.

This company:

- Is looking to cut costs and streamline business operations
- Is looking to obtain a good reward, and is willing to take some risk
- Will just offer its existing products and services, and in its current markets
- Is considering cloud for its complete business process portfolio
- Is looking at a large scope of operation and process complexity
- Wants to improve collaboration between its component companies, and for them to share information (but not for that information to go outside the corporation)

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

3.2.3 ViWi

ViWi is a small start-up company based in California that was formed to develop and sell the product that is its founder's dream – the Virtual Widget. This is an intelligent agent operating on the Internet that people can use to do all sorts of wonderful things. (We don't need to go into details of what these are.) The product will be developed on the cloud, and offered to the public as a cloud service.

This company:

- Wants to make big profits from an innovative product
- Is looking to obtain a high reward, and is willing to take a big risk
- Will offer a new product, in a new market
- Is considering the cloud for software development, and as the platform for its product
- Is looking at a large scope of operation but small process complexity
- Intends its virtual widgets as personal devices, but with ability to share some information publicly

3.3 Assessing cloud suitability – the cloud buyer's decision tree

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.

Is cloud computing right for you? If it is, should you be considering IaaS, PaaS, or SaaS? Should you use public, private, community, or hybrid cloud?

If cloud computing is not right for you, should you consider alternatives? Alternatives that have some of the benefits of cloud computing include:

- Outsourcing all or part of your IT systems management or applications management – you could, for example, outsource data centers, IT management, governance, middleware, or data storage
- Business Process Outsourcing (BPO)
- Using hardware appliances
- Using virtualization

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Outsourcing (including BPO) can deliver some of the commercial and organizational benefits of cloud computing, including reduced capital expenditure (CAPEX) and access to large pools of expertise for specialist tasks. Hardware appliances and virtualization can deliver some of the technical benefits, including the ability to deploy additional resources when needed, and more effective resource utilization.

This section presents a decision tree that helps you to determine whether cloud computing belongs in your vision. It makes the connection between the business use and the enabling cloud technology.

It is an aid to decision-making, not a definitive algorithm. Further considerations, either for or against a cloud solution, may emerge when you develop a detailed architecture.

3.3.1 Assessment considerations

There are ten key questions, whose answers will help you to determine whether to use cloud computing, whether a public or private cloud is appropriate, and whether IaaS, PaaS, or SaaS cloud offerings would best meet your business and technical requirements. They are shown as a complete set in the figure below (Figure 3.5).

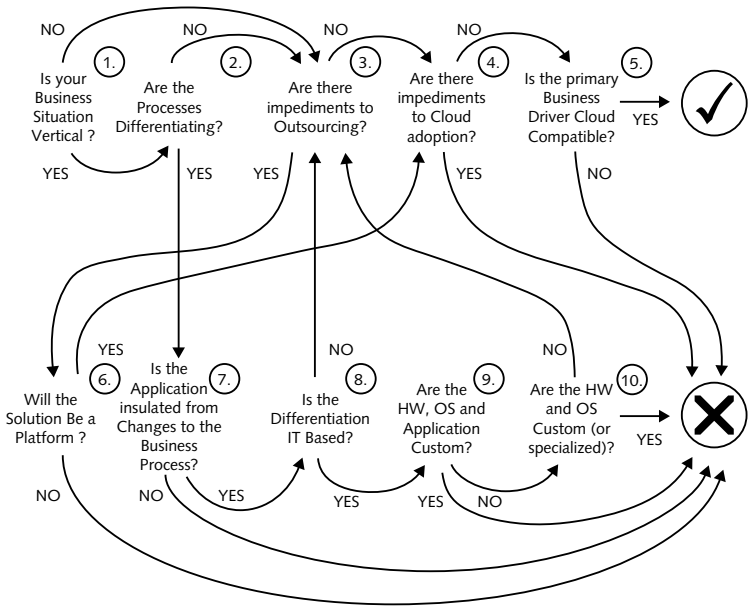

Figure 3.5: Deciding on your cloud solution

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

They are presented as a simple sequence. Your decision process may be more complex and iterative. Answering a question may require you to explore aspects of your business that you had not previously considered, and perhaps to re-evaluate your objectives for using cloud computing. You may find that you need to go back to a question that you answered previously, and re-consider your answer in the light of your later ideas. But thinking about these questions, in the order presented here, will help you to organize your consideration of cloud computing, and lead you to a better solution for your enterprise.

3.3.2 Question 1: Is your business situation vertical?

Is your business scope a vertical or horizontal business process in the industry domain of operation?

Decision	Definition	Next Decision	Objective
Yes: Vertical	You are looking at a single business process or a small number of business processes with consistent requirements that can be considered as a single set. ('Vertical' in this context does not refer to 'industry vertical' or 'geo vertical'.)	Go to Question 2.	Determine which layers of the computing environment shown in the Cloud Service Models figure (Figure 1.3 in Chapter 1) should be outsourced; and – whether outsourced or in-house – for which layers should cloud solutions be considered.
No: Horizontal	You are looking at business processes with multiple sets of requirements that must be considered separately.	Go to Question 3.	Determine how much of each solution layer should be outsourced or in-house, and for which of those decisions should cloud solutions be considered.

Table 3.2: Vertical or horizontal business situation

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

3.3.3 Question 2: Are the processes differentiating?

Is the business process a key differentiator for the business? Are the business drivers and performance of this process critical to the competitiveness and cost-effective operation of the enterprise?

Decision	Definition	Next Decision	Objective
Yes: Differentiating	Processes that are integral to the competitive advantage of the business.	Go to Question 7.	Keep competitors from adopting.
No: Non-differenti-ating	Business-as-usual or external compliance processes that are not significant contributors to competitive advantage.	Go to Question 3.	Keep investments to the minimum required and/ or market parity.

Table 3.3: Differentiating or non-differentiating business process

3.3.4 Question 3: Are there impediments to outsourcing?

Are there any barriers that may prevent some or all of the layers of the solution from being outsourced?

Common impediments to outsourcing include:

- Specification of in-house service cannot be matched externally
- Labor contracts
- Long-term leases
- Switching costs
- Fixed assets with depreciation value
- Immature business architecture
- Business culture
- Geographic location sovereignty rules
- Industry regulation
- Compliance audit rules
- Community relations

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Decision	Definition	Next Decision	Objective
Yes: Impediments to outsourcing	Your business situation includes one or more of the impediments to outsourcing, such as long-term labor contracts and enterprise licenses; aspects of your business environment would impede an outsourcing solution; and these impediments cannot be overcome.	Public cloud is not a good fit. Go to Question 6.	Determine whether a private PaaS (cloud) is a good fit.
No: No impediments to outsourcing	Your business situation does not include any insuperable impediments to outsourcing; your business environment does not have aspects which would impede an outsourcing solution, or the impediments can be overcome.	Go to Question 4.	Avoid pursuing a public cloud or other form of outsourcing solution without first having addressed the impediments to doing so.

Table 3.4: Reviewing impediments to outsourcing

3.3.5 Question 4: Are there impediments to cloud adoption?

Are there barriers to cloud adoption because of issues in hosting the service in a cloud environment that may be detrimental to the business drivers and ROI goals?

Most of the impediments to be considered for outsourcing apply also to cloud adoption. Only additional constraints specific to cloud are listed in this section.

Typical barriers include:

- Custom resources, no major changes planned
- Standard resources, but not adaptable to one-to-many; transformation not practical
- Policy restrictions on resource sharing or on control of configuration changes
- Flat requirements profile, no fractional resources

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

- Too few potential subscribers (not attractive market to providers)
- Entry costs too high
- SLA performance is not acceptable
- The Recovery Point Objective (RPO) and Recovery Time Objective (RTO) performance is not acceptable

These impediments can apply to SaaS, Paas, or IaaS, and to public, private, community, or hybrid cloud solutions.

Decision	Definition	Next Decision	Objective
Yes: Impediments to cloud	Your business situation includes one or more of the impediments to cloud adoption, such as a functionally adequate legacy application on fully depreciated custom hardware; and these impediments cannot be overcome.	Cloud is not a good fit.	Avoid pursuing a cloud solution where this is not appropriate.
No: No impediments to cloud	Your business situation does not include any of the impediments to cloud adoption; or it includes impediments but they can be overcome.	Go to Question 5.	Avoid pursuing a cloud solution without first having addressed any impediments to doing so.

Table 3.5: Reviewing impediments to cloud adoption

3.3.6 Question 5: Is the primary business driver cloud-compatible?

Though there may not be any operational barriers to cloud adoption, if the business objectives are not aligned with the inherent strengths and characteristics of cloud solutions, then non-cloud solutions will be a better fit for the business situation.

Cloud-compatible business drivers may include:

- Reduce medium and/or long-term TCO
- Improve cash flow
- Shift from CAPEX to OPEX
- Improve QoS or SLAs
- Access to functionality or domain expertise

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

- Scale labor or fixed asset capacity up or down
- Become a cloud provider

Cloud-incompatible business drivers may include:
- Cut short-term costs
- Shift from OPEX to CAPEX (public cloud incompatible)
- Increase capacity without need for third-party financing
- Change tax situation (recognize depreciation, job creation incentives, etc.)
- Shift fixed assets (potentially including leases) or labor to provider

Decision	Definition	Next Decision	Objective
Yes: Driver is on compatible list	Cloud strengths match business objectives.	Cloud is a good fit. Review the previously selected Decision Nodes and the Solution Considerations to determine which type of cloud solution is the best fit.	Ensure strategic as well as tactical fit.
No: Driver is on incompatible list	Cloud strengths do not match business objectives.	Cloud is not a good fit. You may wish to consider outsourcing.	Avoid pursuing a suboptimal IT enablement strategy.

Table 3.6: Review the primary business driver

If you determine that cloud computing is a good fit, you should review the following solution considerations.

You should consider all forms of cloud computing if:
- The business process is not a key differentiator for your business.
- The differentiation is not IT-based and the applications are insulated from changes to the business process.
- You have standard hardware, operating system, and applications.

If the business process is a key differentiator for your business, and the differentiation is IT-based, then PaaS or IaaS may fit, and:
- If you have standard hardware and operating system, and custom applications, then SaaS is unlikely to meet your needs.
- If the business process definition is de-coupled from the application, consider SaaS.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

- If the business process definition is not de-coupled from the application, is it time to overhaul the business process, and could the definition be abstracted during that overhaul? If yes, then consider SaaS; if no, then SaaS is not for you.

For example, an SaaS or IaaS solution could be appropriate for a differentiating, IT-supported – or possibly even IT-based – business process if the differentiation is commoditization. There is a well-known PC supplier whose supply chain processes differentiate it from other suppliers. That differentiation is primarily due to the leverage that it has over its suppliers due to the size and frequency of its purchases rather than due to a proprietary supply chain process or IT configuration. Even with the ability to replicate that company's Supply Chain Management (SCM) processes and IT enablement, potential entrants to its market face formidable barriers. So, theoretically, an SaaS solution and perhaps even a BPO solution would be an appropriate choice for enabling that company's SCM processes.

The proliferating online universities in India provide another example of the considerations involved. Should the Indian government force the market to be more efficient by requiring the use of a government-subsidized SaaS or IaaS solution? While consolidation on a cloud solution might be more efficient from a government standpoint, it might not be attractive to the participating universities. But if a university is differentiating its online courses on the basis of content alone, if the courses' resource consumption profile is not complementary to the rest of the university's IT portfolio, and if the government subsidized solution meets that university's QoS and financial requirements, including relief for switching costs, then consolidation would likely be very attractive. So the government must decide whether the consolidation benefits outweigh the costs (including the set-up costs and the unallocated fixed costs prior to reaching break-even), and the risks (including the risk that fewer than the break-even number of universities will make use of the cloud solution).

3.3.7 Question 6: Will the solution be a platform?

Is your objective to transform the solution layers below the business processes and applications into a standard, shared configuration for delivering all of the company's IT services?

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

These layers typically include middleware, operating system, hardware, and data center infrastructure. The middleware may include an Enterprise Services Bus (ESB) to support a Service-Oriented Architecture (SOA).

Decision	Definition	Next Decision	Objective
Yes: Middleware and HW & OS	An IT platform is comprised of the middleware and hardware and operating system solution stack layers. The data center infrastructure may also be included.	Go to Question 4.	Determine whether private PaaS or IaaS plus a platform is a good fit. (Public cloud has already been ruled out.)
No: HW & OS and/ or Data Center	Horizontal HW & OS solutions include server farms, storage pools, and on-demand infrastructure. A data center includes network, physical plant, and a Network Operations Center (NOC).	Cloud is not a good fit, but other solutions such as virtualization, appliances, and out-tasking may be a good fit. Review the Solution Considerations.	Increase the yield of in-house IT infrastructure resources.

Table 3.7: Platform or non-platform solution

If you determine that cloud is not a good fit, you should review the following solution considerations.

If you have standard hardware, operating system, and applications, then any form of outsourcing may meet your needs, and you might use virtualization or hardware appliances.

If you have standard hardware and operating system, but custom applications, then application outsourcing and BPO are unlikely to meet your needs, but outsourcing of data center, IT management, governance, or data storage could be a good fit. Also, you might use virtualization or hardware appliances.

When considering these options, review Question 5 because similar considerations are likely to apply to compatibility of business drivers.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

3.3.8 Question 7: Is the application insulated from changes to the business process?

Can the business process architecture be customized without impacting the underlying standard application configuration and its shared IT management timeline? Can the IT configuration be customized without impacting the execution of the business process with shared resources (especially labor) and using a standard process definition?

Decision	Definition	Next Decision	Objective
Yes: Application is nsulated	Business process definition is abstracted from the enabling application such that a business person with no knowledge of the application can modify the definition of the business process without impacting the ability of the application administrator to manage and maintain the application efficiently.	Go to Question 8.	With insulation, delivery of the IT layers can be somewhat independent of delivery of the business process layer.
No: Application is not insulated	A change to the business process definition requires a change to the application, and *vice versa*.	Outsourcing (including public cloud) is probably not a good fit, and by itself this situation cannot support private cloud. Review the Solution Considerations.	Avoid inhibiting changes to the differentiation; changes to a cloud solution are made at the provider's convenience – whether the provider is internal or external. Insulation and control of change management may not be applicable to a process differentiated purely on scale (i.e., differentiation through commoditization). Private PaaS could be a good fit if the benefits of a common shared IT platform for this differentiating process and several non-differentiating processes, outweighs the benefits of enabling those non-differentiating processes with a public cloud or non-cloud outsourcing solution.

Table 3.8: Reviewing application insulation

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

If you determined that the application is not insulated from the business process, you should review the following solution considerations.

Although it is unlikely that cloud is a good fit, there are some situations where it might be appropriate.

If the application is insulated from the infrastructure, you could consider running it on a public or private cloud IaaS or PaaS.

Even if it is not well insulated, you could consider putting the application on a private IaaS or PaaS if this solution makes sense for your application portfolio as a whole.

To consider these options, go to Question 4.

3.3.9 Question 8: Is the differentiation IT-based?

Is IT is a key enabler of differentiation for the business process?

Decision	Definition	Next Decision	Objective
Yes: Differentiation is IT-based	IT is integral to the differentiation.	Go to Question 9.	Improve IT enablement without enabling competitors to replicate the IT architecture on which the differentiation is based.
No: Differentiation is not IT-based	IT merely supports the differentiating business process.	Business Process Outsourcing (BPO) is not a good fit. Go to Question 3.	Improve IT enablement without enabling competitors to replicate the business process architecture.

Table 3.9: IT-based or non-IT-based differentiation

3.3.10 Question 9: Are the hardware, operating system, and application custom-made?

If the hardware, operating system, and applications layers are all custom-made, then cloud solutions are not a good fit. If only the hardware and operating system have custom elements, or only the applications have custom elements, then cloud might be a good fit.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Mission control weapons systems, and business intelligence systems, are examples of systems where hardware, operating system, and applications layers may all have custom-made elements.

Decision	Definition	Next Decision	Objective
Yes: Custom HW & OS and Application	The IT configuration must be managed on a unique timeline.	Cloud and most forms of outsourcing are not a good fit. See the Solution Considerations.	Legacy IT solutions are often custom, so until a major transformation activity can be justified, cloud is not a good fit.
No: HW & OS and/ or Application are Standard	Either the HW & OS or the Application layer is standard (off-the-shelf).	Go to Question 10.	Investigate the standard IT layer for cloud fit.

Table 3.10: Reviewing custom-made elements in hardware, operating system, and applications

If you determine that cloud computing is not a good fit, you should review the following solution considerations.

If you have custom or specialized hardware or operating system, and custom or specialized applications, then public cloud – whether SaaS, PaaS, or IaaS – is unlikely to be a good fit. Nor are either private IaaS or PaaS, unless there are a number of business processes in the overall portfolio that have complementary resource consumption profiles to this process, that can be enabled with the same IT resources, and cannot be enabled more effectively with alternatives.

Infrastructure outsourcing, application outsourcing, and BPO are unlikely to meet your needs. However, you might consider outsourcing data center operations, IT management, governance, and data storage.

You might also consider hardware appliances and virtualization.

3.3.11 Question 10: Are the hardware and operating system custom-made or specialized?

If these IT layers are standard, you may have a good candidate for a cloud solution.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Examples of custom or specialized hardware and operating systems include those designed for fault tolerance or real-time performance, those including special memory (e.g., supercomputers), and those including cell processors.

Decision	Definition	Next Decision	Objective
Yes: Custom or specialized HW & OS, tailored application	The HW or OS is specialized, while the application configuration fully or nearly matches the most typical configuration of that application.	Cloud and most forms of outsourcing are not a good fit. See the Solution Considerations.	Re-evaluate the requirements necessitating custom HW or OS. If the technology is early stage then, as it matures, look for providers to start to support it.
No: Standard HW & OS, custom application	The HW & OS but not the application are configured in a typical fashion.	Go to Question 3.	Rule out cloud solutions for the application layer.

Table 3.11: Reviewing custom-made elements in hardware and operating system

If you determine that cloud computing is not a good fit, you should review the following solution considerations.

If you have custom or specialized hardware or operating system, but off-the-shelf applications, then no form of cloud computing is likely to be a good fit for your requirements. Nor is BPO or infrastructure outsourcing.

You might, however, consider outsourcing data center operations, IT management, governance, middleware, and data storage. Outsourced application management could work for you if you can use a standard application development platform.

You might also consider hardware appliances and virtualization.

3.4 The example project visions

To illustrate use of the decision tree, let's look at how the example companies would use it to establish and confirm their particular visions.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

3.4.1 Konsort-Prinz

This seasonal novelty company sees that the capacity that it needs to meet its seasonal peak is lying idle for most of the year. Its current systems are old and are not coping with the workload. It needs to replace these systems, and considers cloud computing as an alternative because this provides an opportunity to reduce overall costs by not paying for the unused capacity. Its path through the decision tree is illustrated below (in Figure 3.6), and its answers to the questions are shown in the table that follows.

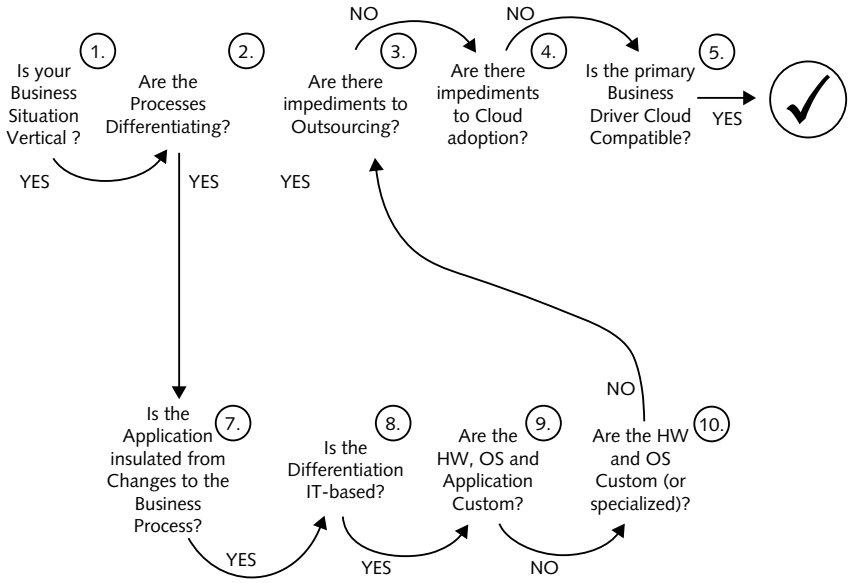

Figure 3.6: Konsort-Prinz decision tree

	Question	Decision
1.	Is your business situation vertical?	Yes, the company is looking at its sales, distribution, and product sourcing processes in combination with enabling applications and supporting infrastructure. This is a small number of business processes with consistent requirements that can be considered as a single set.
2.	Are the processes differentiating?	Yes, it is the efficiency with which it performs these processes that gives the company its market position.
7.	Is the application Insulated from changes to the business process?	This question is harder to answer, because the company has a bespoke stock management application that is a major reason for its success. However, this application is well-designed, and can support many changes to the business process definition without modification. They decide that, on balance, the answer is “Yes”.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

	Question	Decision
8.	Is the differentiation IT-based?	Yes. The stock management application is an important differentiator.
9.	Are the hardware, operating system, and application custom-made?	No, the application is custom-made, but it runs on UNIX on a standard hardware platform.
10.	Are the hardware and operating system custom-made or specialized?	No, the hardware and operating system are standard. Cloud solutions are not appropriate for the application layer, but could be used at the layers below.
3.	Are there impediments to outsourcing?	The company is based in a country that has strong employment-protection legislation, which will make it hard to lay off IT staff. But this obstacle is not insuperable. Decreases in IT staff could be achieved by a combination of incentives to early retirement and generous redundancy payments. So Konsort-Prinz gives answer "No" because the impediments can be overcome.
4.	Are there impediments to cloud adoption?	No, none of the listed impediments apply; there is no reason not to use cloud.
5.	Is the primary business driver cloud-compatible?	Yes, the primary business driver is to scale asset capacity and hence reduce TCO. Cloud is a good fit (assuming that the impediments to outsourcing are overcome).

Table 3.12: Konsort-Prinz decisions

The business process is a key differentiator for the business, and the differentiation is IT-based, so PaaS or IaaS may fit. Konsort-Prinz has standard hardware and operating system, and a custom stock management application, so SaaS is unlikely to meet its needs for stock management, but could be considered for other applications, such as accounting and human resources.

Konsort-Prinz decides to go for public cloud because its primary aim is to allow for large variations in processing load. With private cloud, it would have to invest in sufficient capacity to meet peak seasonal demand, and this would be unused for most of the year. A hybrid cloud, with enough capacity to meet minimum demand in-house might be a possibility, but the application architecture is not suitable for multi-site processing.

The company has established its vision: "We will cope better with peak business demand by deploying our applications software on cloud computing platforms."

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

3.4.2 Sam Pan Engineering

This engineering conglomerate has a different vision. Each of its subsidiary companies has a different IT solution for the same problem. It has a large number of small, but still under-utilized systems, each requiring specialist support. It sees an opportunity to streamline its operations by offering its subsidiaries capacity on better-utilized, more cost-effective systems, by reducing its support burden, and by rationalizing processes through use of common software. Cloud computing could be the means to achieve this.

This company's path through the decision tree is rather simpler.

Its responses to the questions are shown in the table.

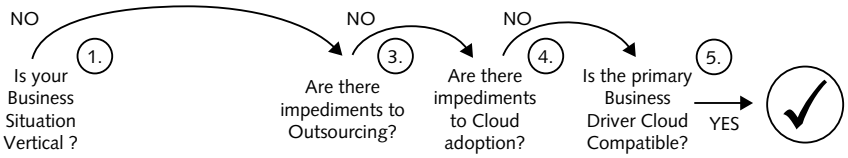

Figure 3.7: Sam Pan Engineering decision tree

	Question	Decision
1.	Is your business situation vertical?	No, Sam Pan has a multitude of business processes with multiple sets of requirements that must be considered separately.
3.	Are there impediments to outsourcing?	No, none of the listed impediments apply.
4.	Are there impediments to cloud adoption?	No, none.
5.	Is the primary business driver cloud-compatible?	Yes. The primary business driver is to increase efficiency. It is certainly cloud-compatible.

Table 3.13: Sam Pan Engineering decisions

Again, the solution considerations indicate that all forms of cloud computing should be considered.

Sam Pan Engineering does not have large variations in workload, and sees regular use of public cloud services as being very expensive. Investment in capacity to support a private cloud looks much more attractive financially. It

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

decides to implement a private cloud, and then to place standard applications on its cloud platform and offer them as SaaS to its constituent companies and departments.

It has established a different vision: "We will reduce costs and streamline our business processes by offering standard software applications on a private cloud."

3.4.3 ViWi

This company has an opportunity to make money from a radically new product. Cloud computing seems to provide an excellent environment for development and launch. It uses the decision tree to validate this. It takes the same path through the tree as Sam Pan Engineering. Its responses are shown in the table.

	Question	Decision
1.	Is your business situation vertical?	No, the processes of developing and operating the virtual widgets have different requirements that must be considered separately. And, in operation, a virtual widget is effectively a platform for certain kinds of business operation, and the business processes for these operations will have multiple sets of requirements that must be considered separately.
3.	Are there impediments to outsourcing?	No, this is a start-up company, and has no impediments to anything.
4.	Are there impediments to cloud adoption?	No.
5.	Is the primary business driver cloud-compatible?	Yes, the primary business driver is simply to make lots of money from a new idea.

Table 3.14: ViWi decisions

Again, all forms of cloud computing should be considered. ViWi goes for public cloud to support its development because this means that it can avoid capital expenditure in the start-up phase. PaaS, or IaaS plus a standard platform, will give the development facilities that it needs. When the product is launched, it will use the same cloud service for the run-time environment. Again, this will mean that it avoids heavy capital expenditure before there are significant revenues to justify it. If the product takes off as hoped, it may invest in some in-house capacity to give a better return.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

ViWi's architecture vision is: "We will develop and launch a new product without the need for major investment in supporting IT by using cloud services."

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Chapter 4

Buying cloud services

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.

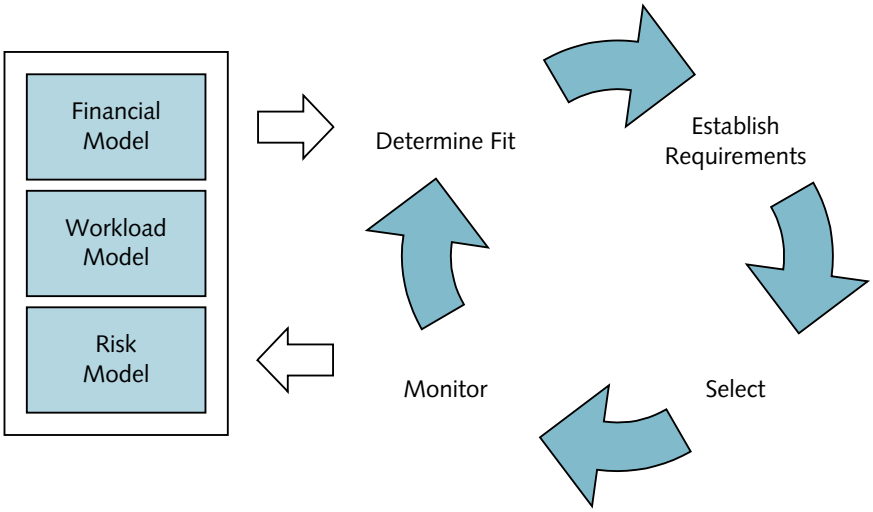

Figure 4.1: The cloud buying lifecycle

The starting point is your system functional requirement plus a set of models of the solution that you are developing.

In the *determine fit* phase, you assess how well the cloud services you are considering fit your functional requirements and solution models. This gives you a first-pass filter on those services. It also gives you some additional solution requirements.

In the *establish requirements* phase, you add to the initial requirements, and establish the complete set of requirements that your candidate services should meet.

In the *select* phase, you select the successful candidate and agree on a contract. The selection depends on which services best meet the requirements, on what they will cost, and on their terms of contract.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

In the *monitor* phase, you observe the performance of the running services, to see how it compares with what was promised or assumed. The models that you started with, and the set of requirements that you established, form the basis for this comparison.

The results of this phase, and the updated models, form the starting point for the next iteration of the cycle, when you renew the service contracts or select alternatives.

4.1 Determining fit

In establishing your cloud vision, you have achieved an understanding of the business context, and made an assessment of cloud suitability. At this point you should have a good idea of whether cloud services will be part of your solution and what kind of cloud services to consider. The next stage is to see whether those services fit your functional requirements and solution models.

For IaaS or PaaS, the functional requirement is simply for infrastructure or platform, unless you need special features – in which case you should probably not be considering a cloud solution – functional fit is not a problem. For SaaS, you must establish that there are cloud services that meet your needs for application processing. For all kinds of cloud service, you then assess fit with the solution models.

Developing the models for your solution will take time and work. Do not expect them to appear magically as a result of your cloud architecture vision. You should develop the models in collaboration with your enterprise architect and the IT team. The enterprise architect will probably take the lead in this. Expect several iterations, with different versions of the models, before you are satisfied with them.

The important models for choosing and monitoring cloud services are the financial, workload, and risk models. The workload model and the financial cost model are intimately involved in the buying lifecycle. The other financial models, and the risk model, are less directly related.

This section describes the workload and cost models, and how they and the other models are used to determine fit. The descriptions are illustrated by the

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Konsort-Prinz and ViWi examples. (Sam Pan Engineering is developing a cloud platform rather than buying cloud services, so is generally not relevant here.)

4.1.1 Workload and cost models

Your workload model is crucial. It enables you to estimate your cloud service costs. It also enables you to determine your throughput and configuration-speed requirements.

The model describes, in business terms, the amount of processing that the service must handle, and how this varies with the number of users and over time. It also describes, in technical terms, the resources that are used to perform this processing, and it shows how the resources used relate to the processing load.

There is no standard workload model. It is different for every enterprise, and for each service that the enterprise uses.

You can estimate the service costs from the resource usage in the workload model. The calculations will unfortunately be different for each supplier, because suppliers package and describe their resources in different ways. This is particularly so for processing resource; the processing unit offered (and charged for) by each supplier may be different, and it can be hard to relate these units to each other. You have to do this, though, if you want to compare the different offerings. This means that you need a cost model for each supplier.

It is a good idea to note how sensitive the models are to changes in different factors. For example, storage costs may be easier to predict than I/O costs. This information will help you to assess risk, as described in the next chapter (Chapter 5).

4.1.2 Workload factors

When modeling workload, you should consider variability and predictability, as well as the average or 'steady state' value.

Average levels might, for example, be 20 transactions per second and 500,000 transactions per day, with variability between zero and 20 transactions per second, depending on the time of day.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Predictability can be:

- Planned and scheduled – for example, Financial Management (FM) mini-peaks at end of quarter and a major peak at the end of the fiscal year
- Planned and unscheduled – for example, pharmaceutical certification on average five times a year, Florida hurricane response
- Unplanned and unscheduled – for example, air traffic control response to ash cloud

Where predictability is planned and scheduled, you should note its periodicity. This is the cycle of requirements, including average, peak, and off-hours. Cycles may be on a calendar basis (including hourly, daily, weekly, monthly, quarterly, and annually), or on an event basis. Examples are: monthly periodicity, with peaks at the same time each month, and annual periodicity, with seasonal peaks.

You should also consider events with predictable capacity, which don't occur at the same time each month or year. Examples are weather patterns such as blizzards, unpredictable fads or events, occasions such as weddings, and disasters such as fires.

Think about the growth or shrinkage of your typical steady state usage, including such things as:

- Increase or decrease in capacity at the beginning of each steady state contract year
- One-time changes
- Acquisitions, divestitures, regulatory changes, and new markets

Consider the average and peak number of users by type. Types of user include business users, 'power' users, and process administrators.

4.1.3 Workload allocations

Be aware that there are different types of workload that scale in different ways. This must be taken into account when modeling your workload.

Applications with a large amount of shared memory, many inter-dependent threads, and tightly-coupled interconnections cannot easily be partitioned between processing units, and only scale vertically. Applications with independent threads that do not share memory and have loosely-coupled

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

interconnections can be partitioned, and scale horizontally. Other applications scale to some extent in both directions.

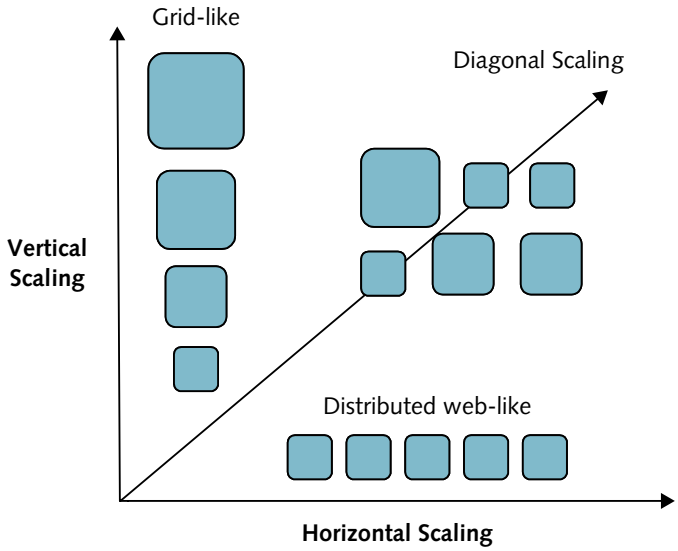

Figure 4.2: Types of workload allocation

The figure below (Figure 4.3) shows some examples of the different types.

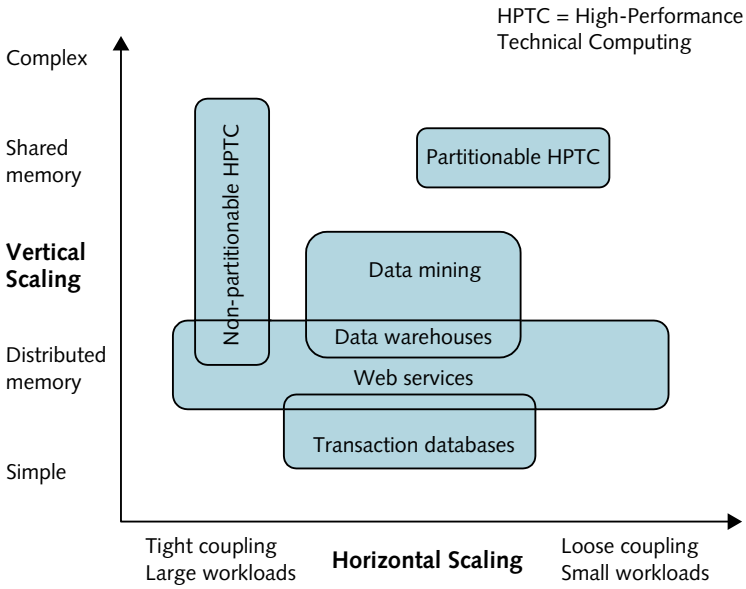

Figure 4.3: Examples of workload allocation types

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Principles to follow in allocating workload include:

- Centralize key services for maximum pooling of total capacity available
- Deploy solutions based on standard commoditized modular building blocks for maximum flexibility and re-use potential
- Automate provisioning of standard images for maximum agility
- Abstract the workload, where appropriate, from infrastructure using virtualization
- Identify workloads that can benefit from moving short-term burst loads to a cloud service

4.1.4 Pay-as-you-go *versus* ownership

Cloud costs are typically lower with short-term burst use but, for a steady or moderately variable load, buying an on-premise system will probably be cheaper than using cloud services in the long term.

There are pricing options that can make cloud attractive in some long-term scenarios. Reserve cloud instances – cloud resources that are contracted for on a longer-term basis and typically for a lower price than those provided on-demand – lower the cost of long-term use. Burst services enable peak operational demand to be diverted to cloud capacity.

Companies should assess their demand usage models in terms of actual workload profiles *versus* capacity usage. Many operational workloads are not 24x7x52 and can be accommodated by a combination of reserve instances and spot instances or 'cloudburst' facilities. The figures below illustrate this. (Note that 30% OPEX is assumed for hybrid cloud and 25% for on-premise. In a pure on-premise solution OPEX is likely to account for 25% of the annual cost with 75% being interest and depreciation – the CAPEX element. The hybrid cloud solution has a lower proportion of CAPEX and a higher proportion of OPEX because of its public cloud element. These amounts are illustrative of what might be encountered.)

With IaaS and PaaS, you can have a hybrid system, with some processing performed on-premise and some on the cloud. (This is not such a practical possibility with SaaS.) The challenge is to define variable workloads and to be able to turn capacity on and off and switch between off and on-premise services. There are costs associated with this that need to be assessed to 'plug into the cloud'.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

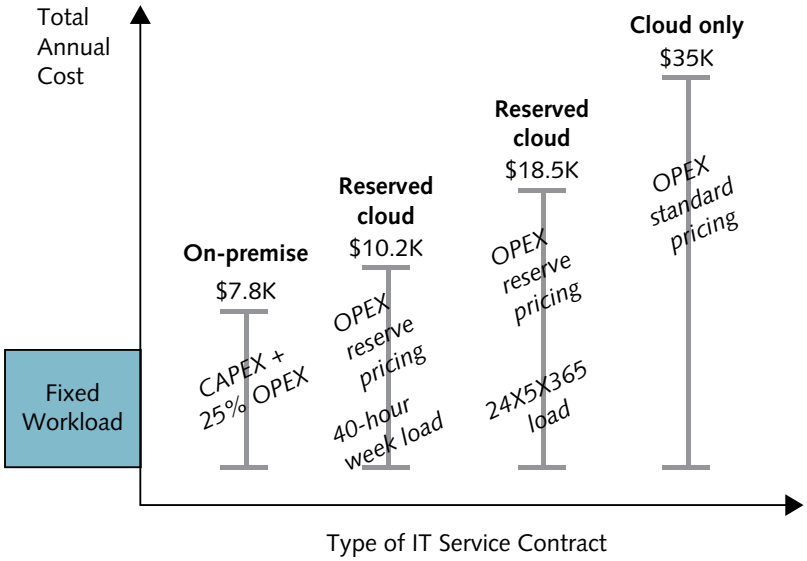

Figure 4.4: Example fixed workload costs

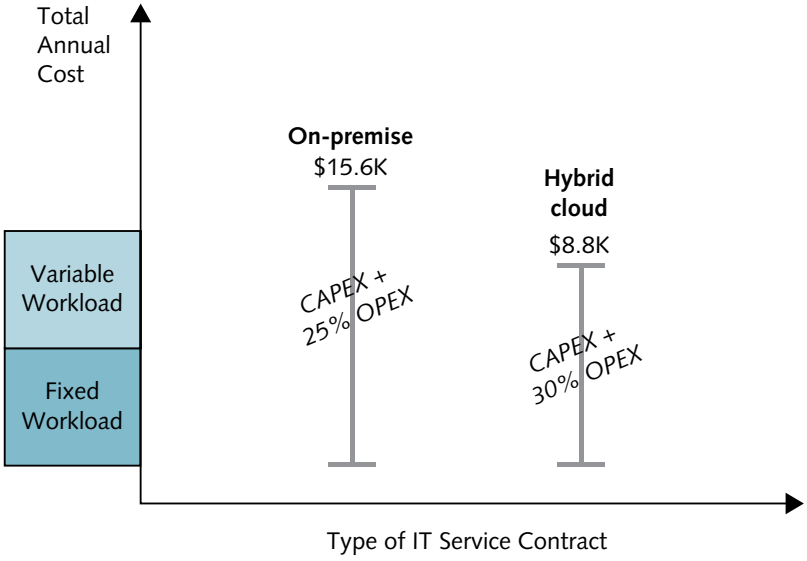

Figure 4.5: Example variable workload costs

4.1.5 Modeling resources and costs

If you are using SaaS, and the charges relate to business units (such as number of users) rather than computing resources (such as Megabytes of memory), there is no need to build the resources into your model. You build a cost model for each supplier by mapping your business units to the supplier's charging units. This model enables you to determine your probable costs for a given load.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

For PaaS and IaaS (or for SaaS if the supplier charges by resource rather than by business unit) you should model use of processing power, memory, storage, and I/O. These are what most suppliers charge for.

Consider:

- Peak resource level – the capacity you require the supplier to be able to provide
- Usage per time period – which is what you will be billed for
- Rate of resource change – which determines your requirement for speed of provisioning and de-provisioning

Try to express these in standard terms that you can map to the units used by the suppliers you are considering. There may be benchmark figures available that can help you. This is a difficult area, which requires specialist technical expertise. If you do not have this in your IT team, consider employing an external consultant.

Mapping your units to the units used by the suppliers will give you the cost models that you need.

It is sensible to validate the models before relying on them for big decisions. You might develop your own benchmark program and run it with test data using each of the services that you are considering. One of the advantages of cloud computing is that on-demand self-service makes it easy to do this.

4.1.6 Example workload and cost models – Konsort-Prinz

Konsort-Prinz develops a simple workload model. The load is measured by the number of products, the number of customers, and the number of orders. The processing power, memory, storage, and input/output (I/O) required are calculated from these factors by simple algebraic formulas.

The company has estimates, based on experience, for typical, maximum, and minimum loads for each month of the year. Feeding these into the model gives estimates for the computing resources required, including the maximum at any time and the expected total for a year.

There are large variations in load, but the levels can be predicted with reasonable accuracy, at least a day in advance. The architects plan to provision

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

capacity twice per day, and set a requirement to be able to provision a new resource within one hour. They are confident that this will meet the needs, and do not see a need to model rate of change.

The modelers express processing resource in 'computer resource' units based on the existing in-house system. They develop a simple benchmark program based on their model, and calibrate it on the in-house system. They then run it on the IaaS services that they are considering, to validate the model and determine how the services' units relate to theirs. This gives them a validated cost model for each IaaS supplier.

The end result is that they can calculate the expected annual resource cost for each supplier, giving a vital part of the information the company needs to make the supplier choice.

4.1.7 Example workload and cost models – ViWi

ViWi's founder draws up a different model. He measures the load in terms of the number of virtual widgets in existence and the proportion of them that are active. Again, the processing power, memory, storage, and I/O required are calculated simply from these factors. Each widget is estimated to require 10 Megabytes of disk storage all the time, and to require 0.1 processing units, 100 Megabytes of memory, and an average of 25 2-kilobyte data transfers per second when it is active.

He has no real idea of how many virtual widgets the company will sell, or what proportion of the time they will be active. For supplier selection purposes, he decides to assume 100,000 widgets active 5% of the time on average and 20% at peak, giving a maximum of 2,000 and average of 500 processing units, with corresponding figures for memory, storage, and I/O.

The predictability is planned and unscheduled. He anticipates that the workload will change dynamically during the day, with sudden, unexpected peaks. The system will provision resources automatically in response to actual and anticipated demand. He sets a requirement to be able to provision or de-provision a new compute unit plus memory within two minutes, and the system will provision processors and memory at 10% above required capacity to allow for unexpected surges.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

He captures the model in a simple spreadsheet.

	A	B	C
1			
2	Storage Megabytes per Widget	10	
3	Processing units per active Widget	0.1	
4	Memory Megabytes per active Widget	100	
5	IO Transfers per second per active Widget	25	
6	Average IO transfer size (Kilobytes)	2	
7			
8	Average proportion of active Widgets	5%	
9	Peak proportion of active Widgets	20%	
10	Spare widget provisioning	10%	
11			
12	Number of Widgets	100,000	
13			
14	Peak processing power	B12*B9*B3	
15	Peak memory usage (Gigabytes)	B12*B9*B4/1000	
16	Peak Storage usage (Gigabytes)	B12*B9*B2/1000	
17	Peak IO transfers per second	B12*B9*B5	
18			
19	Average processing power	B12*B8*B3*(1+B10)	
20	Average memory usage (Gigabytes)	B12*B8*B4*(1+B10)/1000	
21	Average storage usage (Gigabytes)	B12*B2/1000	
22	Average IO transfers per month (Gigabytes)	B12*B8*B5*B6*3600*24*30/10^6	
23			
24			

Figure 4.6: ViWi workload model

The model expresses processing resource in the units used by one particular cloud provider. ViWi's founder (and, at that point, the sole developer) deploys an initial proof-of-concept version of the software on that provider's resources to validate the model, and on other providers' resources for comparison.

The measure derived from the model that he will use to compare suppliers is cost per widget per month.

4.1.8 Using the models

Having created a cost model for each service under consideration, you determine the expected cost of using it in the solution, and feed this cost into the overall financial model. You should also use the risk model to assess the risk profile of the solution, as described in Understanding Cloud Risk (Chapter 5).

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

These financial and risk assessments may result in rejection of some cloud services. Indeed, you may reject all the candidate services, and conclude that cloud computing is not the answer after all.

But if you have made a thorough assessment of cloud suitability it is likely that you will have some candidate services at this stage. You have determined fit, and proceed to establish your requirements.

4.2 Establishing requirements

The requirements are conditions that the solution must or should meet. They can be separated into *mandatory* (must meet) and *optional* (should meet) classes. When you select a service, you reject those that do not meet the mandatory requirements, and give preference to those that meet the optional requirements over those that do not. When you monitor a service in use, you check that it really does meet the mandatory requirements, and how far it meets the optional ones.

(The workload and cost models may be used to determine requirements, and may help you to express requirements, but they are not requirements as such.)

You are likely to have requirements that relate to use and management of services generally, not specifically to cloud services. Standard practice descriptions of this are available from a number of sources, such as the Information Technology Infrastructure Library [ITIL]. You may wish to refer to one of these descriptions when establishing your complete list of requirements. This section focuses on requirements that are specific to, or affected by, cloud computing.

You will want to describe requirements for functionality, supplier choice, performance, manageability, security, and regulatory compliance. The requirements in these areas that are described in this section are summarized in the table.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Area	Requirements
Functionality	Service functionality Back-up Bulk data transfer
Supplier choice	Supplier choice
Performance	Availability Reliability Recoverability Responsiveness Throughput
Manageability	Configurability Reporting Fault management
Security	End user access control Provider access control Resource partitioning Logging Threat management
Compliance	Compliance with regulations

Table 4.1: Cloud computing requirements areas

Different requirement areas may have greater or lesser importance depending on whether you are considering IaaS, PaaS, or SaaS.

4.2.1 Service functionality

These are the requirements for what the system should do. They are considered as overall solution requirements when establishing the suitability of a cloud solution, as described under Determining Fit (Section 4.1), and they then translate directly to requirements on the cloud services that will be the basis of that solution.

IaaS enables the consumer to provision processing, storage, networks, and other fundamental computing resources. You should consider what resources you need. In particular, you should consider what operating system you will use, and whether you will obtain it from the provider.

Konsort-Prinz needs UNIX (or Linux) to run its applications, and wishes to obtain it as part of the cloud service rather than loading it onto bare computing resources. The company notes this as a requirement.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Sam Pan Engineering is in the position of providing a cloud service, rather than consuming one. Its users have a mixture of UNIX and Windows applications, and the company notes that the ability to support both platforms is a requirement on the cloud service that it provides.

PaaS enables the consumer to deploy applications created using programming languages and tools supported by the provider. You should specify both the environment in which the applications will run, and that in which they will be created.

One possibility is that much of the development is done on local machines, and the code is only uploaded to the cloud service for integration, final testing, and deployment. This saves cost, and can dramatically speed up development and unit testing. You should consider this and, if appropriate, specify requirements for it.

ViWi is writing its virtual widget program in Java. It requires a Java programming environment and a Java run-time environment. The founder plans to do most of the development on local machines, using the Eclipse [ECLIPSE] open-source software development environment, and make the ability to support this a requirement also.

SaaS enables the consumer to use the provider's applications running on a cloud infrastructure. You should specify the functional requirements for such applications just as you would if you were planning to run them in-house.

4.2.2 Back-up

Regular back-ups enable data to be recovered in the event of system failure. They can also be useful when users wish to correct mistakes that they have made.

The provider may take responsibility for backing up data. In fact, on-premise cloud provider involvement is unavoidable when, for example, tape back-up is being used. You may wish to specify recoverability objectives (see the section on Recoverability below) rather than particular back-up patterns.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

In some cases, the back-ups will be your or your users' responsibility. You should specify the features that the service must have to enable these back-ups to be taken and the data to be restored from them.

Care should be taken that the back-up and restore do not disrupt normal operations.

Konsort-Prinz plans to take a daily back-up, to match the current practice with its in-house systems. It specifies a requirement for their provider to do this, and to keep the back-up copies in safe storage.

4.2.3 Bulk data transfer

10 Megabits per second sounds like a pretty good data transfer rate on today's Internet. But it will take nearly a quarter of an hour to transfer a Gigabyte at that speed; or about 10 days to transfer a Terabyte. We still have to worry about bandwidth when moving the data volumes that are common in the modern enterprise.

You may need to transfer large volumes of data in the normal course of business. You should also think about the need to transfer all of your data if you move to another supplier, as part of your exit strategy.

Some cloud suppliers provide bulk data transfer facilities; for example, using physical transfers on disk packs.

You should consider your possible bulk data transfer needs, and specify facilities to make them feasible.

4.2.4 Supplier choice

Making it a requirement that there should be other similar suppliers that you could move to if necessary will enable you to have a viable exit strategy. The importance of this is discussed under Selection (in Section 4.3). It will strengthen your hand considerably in negotiations.

On top of TCP/IP, the cloud uses established standard web and web service data formats and protocols. The programming interface standards on which cloud PaaS offerings are based are equally well established. They include the single-vendor .NET standards and multi-vendor standards such as

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

UNIX, Linux, and Java. As regards basic service functionality, this ensures a reasonable choice of supplier for IaaS and, in theory, for PaaS – though there are not yet all that many PaaS suppliers to choose from.

The situation is different when it comes to configuration and management, which can vary substantially between cloud suppliers, even for IaaS. It is unlikely that standards can make it possible for an enterprise to have a single management regime for all its cloud suppliers, but they should at least make it easier to move from one supplier to another. At present, however, this is academic. Widely accepted standards for cloud service configuration and management do not yet exist. There are several industry bodies working on them, including the Distributed Management Task Force [DMTF], the Open Grid Forum [OGF], and the Storage Networking Industry Association [SNIA].

4.2.5 Availability

Availability, or *uptime,* is the proportion of the time that a system is available for use. It is commonly described by the number of 9s. A 'Five 9s' system is up 99.999% of the time – a little over five minutes per year downtime. Planned, scheduled outages for maintenance are typically excluded.

Availability is important. An uptime of 99.9% means 42 minutes of downtime per month during which you cannot provide service to your customers.

The Uptime Institute [UPTIME] defines a widely used standard set of availability tiers for data centers, ranging from 99.999% at Tier 4 down to 97.6% availability at Tier 1 (nearly 9 days downtime per year). These levels can be used as a yardstick for cloud service availability. Historically, the better data center service levels (four or five 9s) have not been matched by cloud service providers, but this is changing as reliability and commercial models evolve and the providers develop their recovery and performance management mechanisms.

ViWi wants availability to be in the upper range of what cloud services typically offer. It specifies an availability requirement of 99.9%.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

4.2.6 Reliability

In the world of electronic components, reliability of repairable components is expressed as a combination of two parameters: Mean Time Between Failures (MTBF) and Mean Time To Repair (MTTR). For non-repairable components, it is expressed as Mean Time To Fail (MTTF). By analogy MTBF and MTTR can be, and often are, used as a measure of reliability of cloud services.

There is a simple algebraic relation between average availability, MTBF and MTTR [average availability = MTBF/(MTBF + MTTR)]. You would probably not specify reliability as well as availability. You would specify reliability rather than availability if one of the two reliability factors is more important to you than the other.

MTBF and MTTR are average figures. They are useful as indicators of the expected quality of the service, but they are not useful in the context of an SLA. Suppose the MTTR of a service is quoted as 30 minutes and you have an outage of one hour. Is the provider in breach of the agreement, or were you unlucky in suffering a particularly long outage, balanced in the average by shorter ones at other times? In an SLA, you should look for guaranteed maximum number of outages over a period and maximum outage length, and similarly for a guaranteed uptime during any month or year. These figures will of course not be so good as the averages for the service.

Konsort-Prinz believes that its customers will accept failures as part of a normal Internet experience, but that the customers will become frustrated, and perhaps turn to other suppliers, if the system is down for a long time. It therefore puts a higher priority on MTTR than on MTBF, and specifies a maximum outage length.

Sam Pan Engineering, on the other hand, will re-allocate users to another resource if one of its cloud resources goes down, but wants to do this as seldom as possible. Its priority is on MTBF rather than MTTR, and it specifies maximum number of failures per year.

4.2.7 Recoverability

Recoverability is the ability to recover from a failure. It is measured in terms of Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

RTO determines how quickly the system needs to be fully operational; RPO determines how much data loss can be tolerated.

RTO and RPO are normally used in the context of catastrophic failures, where there is a possibility of a large amount of data being lost (e.g., the contents of a whole disk), rather than transient failures (e.g., an operating system crash) where just a few records are likely to be lost. These situations should be distinguished; they require different kinds of recovery strategy.

RTO is specified as a time period and includes the time taken to restore data from back-up copies.

RPO can be specified in terms of the number of hours' data that can be lost. In some applications, such as financial transaction processing, this may be zero – but satisfying such a requirement is likely to be expensive. The daily back-up specified by Konsort-Prinz implies an RPO of 24 hours' data.

4.2.8 Responsiveness

Response time is the time taken for the system to respond to input. It can be specified as an average with an allowance for variability; for example, "an average of 300 ms with no more than 1% of responses taking more than 800 ms". Alternatively, a maximum figure may be given; for example, "less than a second".

When considering the overall system, different response times may be appropriate for different kinds of interaction. For example:

- Product web page displayed in less than two seconds
- Web purchase completed in less than 30 seconds

The time to respond to the end user is the most important response time figure. Unfortunately, you will never be able to specify that to a cloud supplier, because it always includes factors that are outside the supplier's control. In all cases there are network delays between the cloud system and the end user. For IaaS and PaaS there are also delays due to application software and perhaps other programs that are your responsibility.

The most that you can specify is the processing and communications latency within the cloud system itself. This can be quite low; figures of the order of 1

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

millisecond (ms) are quoted by some providers. In contrast, for exchange of information over the Internet between a user and the service – which is what most consumers care about – you can expect a latency of 100 ms or more, perhaps even a second or more.

Part of the round-trip delay may be due to slow response by the provider. Cloud system response should be broadly comparable with that found in a non-cloud system. However, higher system utilization generally means slower response, so you should expect the response of a highly utilized cloud system to be slower than that of an under-utilized system that you have in-house.

Also, there are sometimes additional delays due to virtualization. For example, there is a particular consequence of some resource pooling arrangements that can be a problem – known as 'cold start' – on cloud systems. When a virtual processor has not been in use for a certain period, it is removed from its real processor to make room for others. When it is next required, it has to be re-loaded. An application that runs infrequently may thus have to be re-loaded every time, so that the response perceived by the user is very poor indeed. This has little effect on the response time averaged over all applications (which is probably the figure that the supplier will quote) but a big effect on one particular application (which may be the one that matters to you).

The lesson is that it is advisable to understand the cloud supplier's architecture, and the way that your own software works, when specifying response time requirements.

Konsort-Prinz does not have unusual response time requirements, and believes that its application software will run on the cloud systems with similar response times to those that it has when running in-house. It does not specify a particular response time requirement for the cloud services that it will use. [Its management team is in for a nasty surprise – see Measuring and Tracking ROI (Section 6.6).]

ViWi does not have unusual response requirements for user interactions, but it does have a need for frequent, rapid communications between the virtual widgets and a central support processing unit. It specifies latency of maximum 1 ms for these communications.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

4.2.9 Throughput

Throughput is the amount of work that a computer can do in a given time period. For transaction processing systems, it is normally measured as transactions-per-second. For systems processing bulk data, such as audio or video servers, it is measured as a data rate (e.g., Megabytes per second). Web server throughput is often expressed as the number of supported users – though clearly this depends on the level of user activity, which is difficult to measure consistently.

The size of each transaction can be important. A system that can process 20 transactions per second, each of which uses a Megabyte of memory, may not be able to process 20 (or even 2) transactions per second if each transaction uses a Gigabyte. Other factors can also be important. The Transaction Processing Performance Council [TPC] defines several different kinds of transaction, with different characteristics, for benchmarking purposes.

The requirements for peak and average throughput should be derived from the workload model, as described above under Modeling Resources and Costs (in Section 4.1).

4.2.10 Configurability

On-demand self-service is an essential characteristic of cloud computing. Consumers can provision capabilities without requiring human interaction with the service provider. But the facilities provided for this vary between providers. You should consider what facilities you need, and specify them as requirements. These might cover:

- The ability for configuration and provisioning to be done by software, rather than manually
- The speed of provisioning and de-provisioning resources
- The ability to specify particular resource combinations
- The ability to set limits on the amount of resource used
- The visibility and controllability of multi-tenancy partitioning and isolation

For example, ViWi needs to provision and de-provision software resources rapidly and under software control, whereas Konsort-Prinz can accept slower provisioning and de-provisioning, and can do it manually.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

4.2.11 Reporting

The reports on system usage and performance must give sufficient information for your purposes.

Reports on usage of billed resources are normal, and form the basis for billing. You may, however, have requirements for additional information to be shown. For example, you might want to know the CPU and memory utilization of the processing units you have deployed. This is the case, for example, for ViWi, whose system uses the information in tuning the allocation of virtual widgets to processing units.

You may also have requirements for the timeliness of these reports. Konsort-Prinz wants the usage reports for each day to be available by 6:00 am the next morning.

Finally, you may want the reports to be in a particular format, or to be in a form suitable for machine processing. Konsort-Prinz wants them to be clear and easy to read by humans. ViWi wants them in machine-readable XML for software analysis.

4.2.12 Fault management

Diagnosis and correction of faults in the service will be the provider's responsibility. You should require adequate procedures to enable you to report faults and check what progress has been made in fixing them. You may require web or telephone help desk and support to be available.

4.2.13 End user access control

Where cloud services are simply used to replace a component of a conventional IT architecture, as is the case with Konsort-Prinz, the end user access control requirements are similar to those of the conventional architecture. But where cloud computing is leveraged to provide a distributed collaborative environment, a whole host of new issues comes into play.

If Konsort-Prinz starts to move beyond its original concept, and thinks about collaborative working with and between its suppliers, it will face some of these issues. What level of collaboration is needed for the supply chain and channels? What information and services must be collaborative and shared?

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Once it has decided this, it must address the issue of how to identify the people who will use its system, and how to ensure that those people can access the data, and only the data, that they should be able to access. The traditional approach is to assign everyone a Konsort-Prinz ID and password, but the limitations of this have been apparent for a long time now, and new paradigms of federated identity and access control are emerging.

Selecting an appropriate mechanism for identity and access control for cloud services is still something of a challenge. Organizations that can rise to it will make support for the appropriate protocols a requirement. For SaaS, this means requiring that the services support the protocols. For PaaS, it means requiring that the platform enables you to build and deploy applications that support the protocols. For IaaS, it means finding a standard platform that does this and will run on the IaaS infrastructure.

ViWi is facing these issues already. Its virtual widgets access third-party applications on their owners' behalf. ViWi intends to use the OAuth protocol [OAuth] for authorization, and makes support for OAuth a requirement for the PaaS service that it is looking for.

4.2.14 Provider access control

The ability of system administrators and other staff with special privileges to access restricted information, even when there is no business reason for them to do so, has long been an issue. Cloud computing gives it a new dimension. An in-house system administrator is a staff member, subject to company checks and discipline. The administrator of a public cloud service works for another company and in a remote location where you cannot see what he or she is doing.

This is a problem recognized by cloud suppliers, and some of them have mechanisms in place to control the access that their staff have to customer information.

If security of your data is important to you, you should make provider access control a requirement, and insist that it is a condition of your contract.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

4.2.15 Resource partitioning

An implication of resource pooling on public cloud is that your applications and data may be sharing resources with other programs that could have been written and deployed by anyone. There is nothing to stop a hacker writing a program and running it on a PaaS or IaaS cloud just as you do – and that program could be running on the same physical CPU as yours, and sharing the same memory.

Cloud suppliers do, of course, have mechanisms to ensure that the programs are securely partitioned. But you might still want to make security from access through shared resources a requirement, and a condition of your contract.

You might also want to have some visibility of and control over what other programs are sharing resources with you. You might make this a requirement, but insisting on it will probably restrict your choice.

4.2.16 Logging

A log of user activity is extremely valuable in the event of a security breach, and can also be important for normal activity. It provides an audit trail that helps the system manager to establish what damage has been done, re-secure the system, and take measures to prevent future breaches of the same kind. Security logs have long been common practice. They may be required to comply with legislation, but you may well want one in any case to meet a business need to keep track of user access. You should specify this as a requirement.

4.2.17 Threat management

Threat management encompasses the disciplines of event, incident, and problem management as defined by ITIL. These are collectively grouped under the term 'threat management' for the purposes of our discussion.

You should look at what procedures the service provider has for evaluating and responding to identified security threats. If a security breach does occur, much of the remedial and corrective work must be done by the service provider. You should make it a requirement that the supplier has procedures that meet your needs.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

4.2.18 Compliance with regulations

Many legal regulations preclude using the cloud as-is. For example, the EU Data Protection Directive 95/46/EC imposes restrictions on where personal data can be held. Apart from the legal regulations, contractual or moral obligations may require you to take care that information is kept confidential, or to guarantee that it is not lost or destroyed. Can you meet these obligations when the data is on the cloud?

You should identify the obligations that impact on your ability to use cloud services. The following are examples that apply in certain jurisdictions. Similar legislation may apply in other jurisdictions.

- The EU Data Protection Directive 95/46/EC [EU 95/46/EC], which applies to the processing and movement of personal data
- The US Sarbanes-Oxley Act of 2002 (SOX) [SOX], which was enacted as a reaction to a number of major corporate and accounting scandals such as that affecting Enron, and sets standards for corporate auditing, accountability, and responsibility
- The US Federal Information Security Management Act of 2002 [FISMA], which requires federal agencies to implement information security programs
- The US Gramm-Leach-Bliley Act [GLBA], which provides limited privacy protections against the sale of private financial information
- The US Health Insurance Portability and Accountability Act of 1996 [HIPAA], which includes provisions that address the security and privacy of health data
- The International Standards for Assurance Engagements (ISAE) No. 3402 [ISAE 3402] or the Statement on Standards for Attestation Engagements (SSAE) No. 16 [SSAE 16], which mirrors and complies with ISAE 3402 and effectively replaces the Statement on Auditing Standard 70 (SAS 70) Type II; these standards apply to assessment and audit of the internal controls of service organizations
- The US Department of Defense (DoD) Information Assurance Certification and Accreditation Program [DIACAP], which defines the risk management process for DoD information systems
- The Payment Card Industry Data Security Standard [PCI-DSS], which states international security requirements for protecting cardholder data

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

For each of the obligations that impact on the cloud services that you use, specify requirements on those services that will ensure that you comply.

Unfortunately, the industry does not in general have adequate certification schemes that assure compliance and QoS delivery of cloud services to customers. In addition, most providers will not allow external audits to be conducted. This leaves the customer in a difficult position, especially towards external regulators. The Open Group has done significant work on the automation of regulatory compliance, resulting in the publication of the Automated Compliance Expert Markup Language standard [O-ACEML], but it has no current plans for a certification program.

Konsort-Prinz will hold some personal data on its system. To comply with EU legislation, the company makes it a requirement that the resources and networks used by the IaaS service that it uses should be located in the EU.

4.3 Selection

Once you have determined the costs of using the services, and established the requirements, you can review contract and delivery terms, negotiate to try to reduce price or improve performance, and select your preferred service. But, before doing these things, consider the question of exit strategy.

4.3.1 Exit strategy

Vendor lock-in has been an issue for customers of IT systems since the early days of computing. Developing open, vendor-neutral IT standards and certifications to enable interoperability and prevent lock-in is the primary role of The Open Group.

The battleground has changed over the years. At first, the concern was with operating systems and applications portability, then with networks and interoperability. Now that we have a reasonable degree of portability, and the Internet has solved the problem of network interoperability, The Open Group is working to improve access to integrated information, within and among enterprises: its vision of Boundaryless Information Flow.

Cloud computing is opening up a new front. The protocol, data format, and program interface standards for using cloud services are mostly already in

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

place. This is why the market has been able to grow so fast. But standards for configuration and management of cloud services are not yet there. And the crucial contextual standards for practices, methods, and conceptual frameworks are still evolving.

Cloud computing will not reach its full potential until the management and contextual standards are fully developed and stable. While these standards are being developed, enterprises will find it harder to use the cloud than it should be. In particular, they will find it harder to change cloud suppliers. It is tempting to call for instant solutions, but history shows that effective standards can only be created in the light of experience.

Today it may be easy to switch mobile phone providers, but the ease of switching cloud providers is still largely untried. As the market becomes more competitive, we would expect the transition between service providers to become easier.

When negotiating with a new supplier, or re-negotiating with an existing one, or if you wish to in-source a particular service, you should be sure that you have an exit strategy – a way of moving to another supplier if you wish to do so.

Insisting on requirements for supplier choice and bulk data transfer will help you achieve this. You should also consider the contract terms (discussed in the next section) and the cost of making the change. You will have to pay for any bulk data transfer – different suppliers charge different amounts. You may also need to convert the data to different formats, re-write software, re-test your solution, and re-educate users.

The problem is relatively straightforward for IaaS and PaaS, because the processing logic is supplied by the user, and the data is likely to be stored in a standard form. It is wise to check, though, that the particular standard formats that you use can be supported by other cloud providers, or that there is a transformation mechanism. For example, if you use SQL, can you easily migrate to a cloud provider that supports persistent data objects? A project that faced this problem recently found that it took two months of solid programming effort to make the change.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

The problem is typically harder for SaaS, because the processing logic is supplied by the provider, and the data formats may be proprietary. (This is similar to the situation with a purchased application deployed in-house.) It is unlikely that another provider will have the same processing logic, and a change of provider will therefore probably mean changes to the business processes. Custom code may be needed for data transformation.

4.3.2 Contract terms

Contracts cover a wide range of aspects, and may be complex. See Legal and Quasi-Legal Issues in Cloud Computing Contracts [MCDONALD] for an overview. This book does not aim to be a comprehensive guide to them, or to give legal advice.

Here are some key questions that relate specifically to cloud computing that you should ask when reviewing cloud service contracts:

- Who owns the data and, for IaaS and PaaS, who owns the programs? What rights does the service provider have over them?
- How long does the contract last, and what possibilities do you have for early termination? This is important for your exit strategy.
- What possibilities does the supplier have for early termination? This affects the risk that you take.
- Does the contract introduce dependencies so that, for example, you can only use the service if you are also subscribed to a different service? This is important for your exit strategy too.
- What is the SLA? This defines the service levels, particularly for performance, that the supplier contracts to support. Typically, it will not cover all of your performance requirements.
- What happens if the supplier does not meet contracted service levels? Some allowance for downtime is routinely made in the contracts with cloud providers, but it is difficult to reclaim losses over and above a refund of the fees charged by the provider. You may want to be able to terminate the contract if the supplier persistently fails to perform.
- Where will the data be located? Can you ensure that it is located in a particular place? You may need to do so to comply with regulations that affect you.
- What security mechanisms does the supplier implement? Particular mechanisms may be needed to comply with regulation.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

- What jurisdiction applies? If you have to take the supplier to court, will you have to do it in a foreign country?
- Can the supplier subcontract? If so, do the responsibilities – for example, for meeting service levels – remain with the supplier, or pass to the subcontractor?

4.3.3 Negotiation

It is not usual to modify standard contractual arrangements, especially if the purchasing company is small and the supplier is large, or if cloud computing is being purchased by the business users and not by the IT department.

Konsort-Prinz feels that it may be able to modify the contract terms of the smaller IaaS providers that it is considering. ViWi sees no chance of doing this with its potential providers (at least until virtual widgets become wildly successful).

If the functional characteristics of your workload are complementary to the other workloads enabled by a provider (i.e., if your workload improves the provider's diversity factor), then your business will be more attractive to that provider, and you may be able to negotiate better terms and conditions. (Konsort-Prinz is unfortunately not in this position, as its peak is likely to coincide with high traffic from other service consumers.)

Even though you cannot modify the standard conditions, it is often possible to select the best combination from a range of options published by the seller.

Duration of contract is a particular example. Traditionally, this might be a point for discussion during negotiations. A cloud supplier may not be prepared to discuss it, but may offer a range of different packages in which the buyer commits to the service for different lengths of time, and pays different prices. Capability can be purchased on demand, but 'reserved' resources cost less than 'spot' purchases, and the longer you reserve them for, the less they cost.

4.3.4 Choosing the service

The choice of supplier is made by trading off the factors of cost, risk, ability to meet requirements, exit strategy, and terms of contract.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Cost is determined from supplier cost models, as described earlier in this chapter, and risk is assessed through the risk models described in the next chapter.

You should prioritize your requirements, identifying those that must be met and those that you are willing to drop in order to reduce cost or risk. For example, you may decide that the peak availability and throughput requirements must be met, but the response time requirements may be diminished. You may have to make similar compromises on exit strategy or contract terms.

Konsort-Prinz has a good choice of IaaS suppliers, including large international ones and some small local ones. It can negotiate better terms with the smaller suppliers, but assesses these as being of higher risk. It has no serious problems with exit strategy; the application runs on a standard platform, the data will not need conversion, and its volume is not enormous. There are no major difficulties with contract conditions. It decides to partner with a regional supplier that offers good terms, accepting the higher risk as compared with a large international supplier.

ViWi is in a very different position. It cannot find a single PaaS supplier that meets all of its specialized and stringent requirements. There is one (which it is already using for initial development) that comes close. The founder decides to consider IaaS as an alternative. He finds an IaaS service that, with a standard platform, meets more of the requirements than most of the PaaS suppliers.

He develops cost models for two PaaS providers and one IaaS provider. These show similar figures, with a reasonable cost per widget per month in each case. (The example figures have been taken from the public websites of real cloud providers, but the names are omitted as the authors do not wish to advertise particular companies, or to risk people forming judgments on them from isolated examples of their pricing structures.)

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Costs per Month	PaaS 1	PaaS 2	IaaS
Price per processing unit per month	$72	$86	$68
Price per storage Gigabyte per month	$0.15	$0.20	$0.20
Price per Gigabyte IO	$0.11	$0.13	$0.13
Processing cost per month	39,600	$47,520	$37,620
Storage cost per month	$150	$200	$200
IO cost per month	$71,280	$84,240	$84,240
Cost per widget per month	$1.11	$1.32	$1.22

Figure 4.7: ViWi cost models

The first PaaS provider meets the requirements best, and comes out best for cost too. The founder is concerned at not having a good exit strategy, but there is little that he can do about this. He chooses the first PaaS supplier, and makes contingency plans for an exit strategy in which ViWi would use the IaaS supplier's infrastructure and a standard platform that is as similar to their PaaS supplier's platform as possible. This would, however, still mean significant software re-development costs.

4.4 Monitoring

While the service is in operation, you should monitor how well it agrees with your workload and cost models, and how well it conforms to the requirements. The results will help you assess whether to renew the service contract, renegotiate a different contract with your supplier, or change to a different supplier. They will also feed in to your ROI and risk assessments, as described in the next two chapters.

4.4.1 Workload and cost

By monitoring workload and used cloud capacity you can see how well your model is working, and keep control of your costs.

Konsort-Prinz keeps track of used and provisioned processing power. This shows that the provisioned processing power has, as planned, been much lower for most of the year than what was used on the in-house system.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

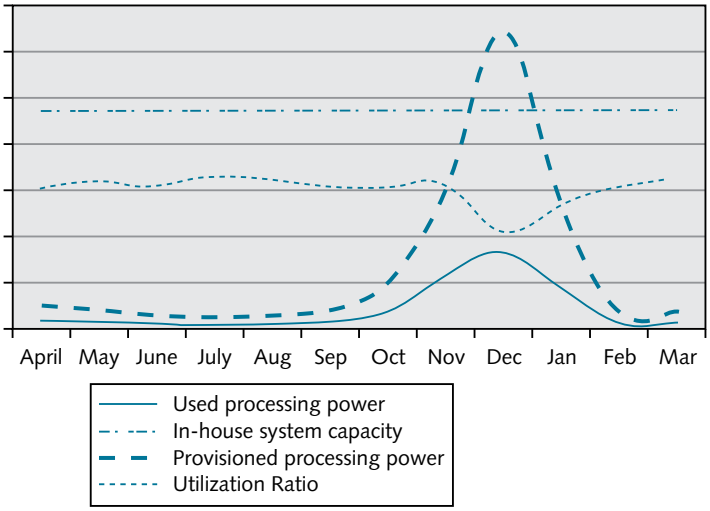

Figure 4.8: Konsort-Prinz capacity utilization

It is still, however, higher than necessary. Utilization of the purchased cloud capacity was only about 33%. The figure below (Figure 4.9), showing hourly utilization for the peak day in December, explains why.

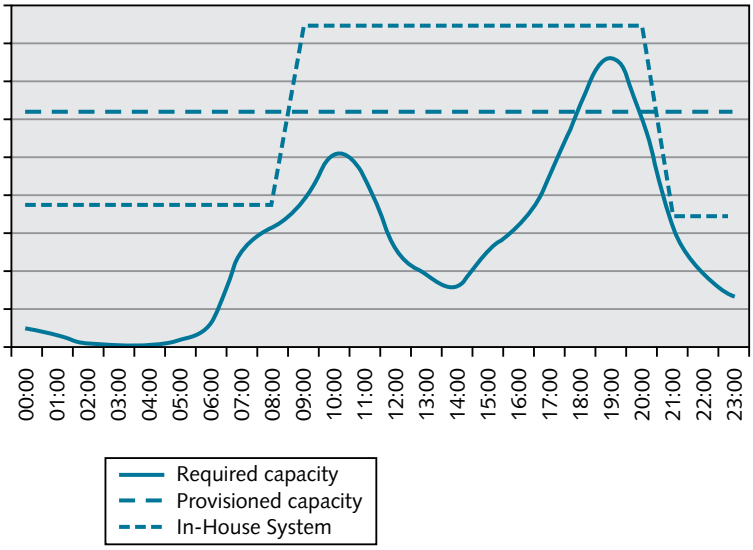

Figure 4.9: Konsort-Prinz peak day utilization

The twice-daily provisioning cycle does not have sufficient granularity. It is just providing sufficient capacity in the on-peak period and in the off-peak period, but this is still much more than is needed at the troughs

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

of the daily cycle. It looks as though there is scope to improve the return further, by fine-tuning the provisioning. This would, however, be difficult, given Konsort-Prinz's 'vertical scaled' architecture. Changing the system configuration is rather an upheaval, and even the twice-daily cycle is proving to be a strain. With a more modern software architecture, based on loosely coupled services that could be separately provisioned with resource, the company would have been in a better position to take full advantage of the utilization improvement offered by the cloud.

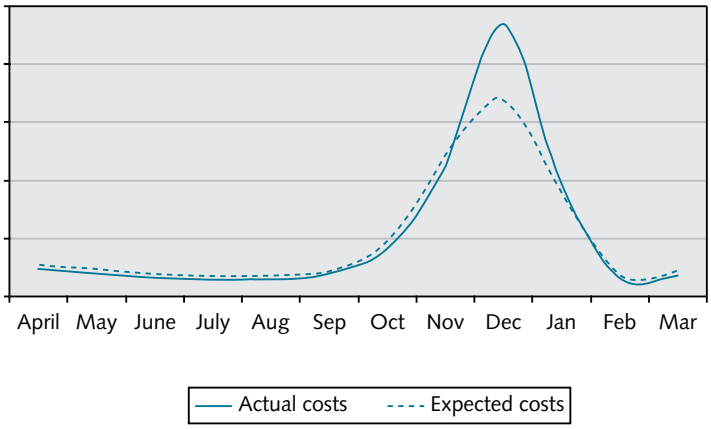

Figure 4.10: Konsort-Prinz first year costs

It monitors actual and projected costs over the first year of operation.

This shows a very good fit of actual to projected costs. The discrepancy in December is largely due to inability to match capacity to load quickly enough in a period of rapid change. [There is a story behind this, which will be told in Understanding Cloud Risk (Chapter 5).] Despite this discrepancy, the fit to the model is close enough. The company executives consider that the new system has pretty much met their cost targets, and they are happy with their supplier on this score.

4.4.2 Conformance to requirements

In many cases, conformance to requirements is "yes" or "no". The service either does or does not perform the required functions, does or does not support the specified protocols, and so on. But for performance requirements, and some others, it is less clear. Suppose that availability was specified at 99.9% and was measured over the year at 99.7%, for example. This may not represent a serious discrepancy in most business environments.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

One approach is to rate the values on a simple scale, from unacceptable to excellent, and mark them up on a scorecard. The figure below (Figure 4.11) shows a comparison of performance and manageability parameters for the Konsort-Prinz cloud-based system with its previous system, using this approach.

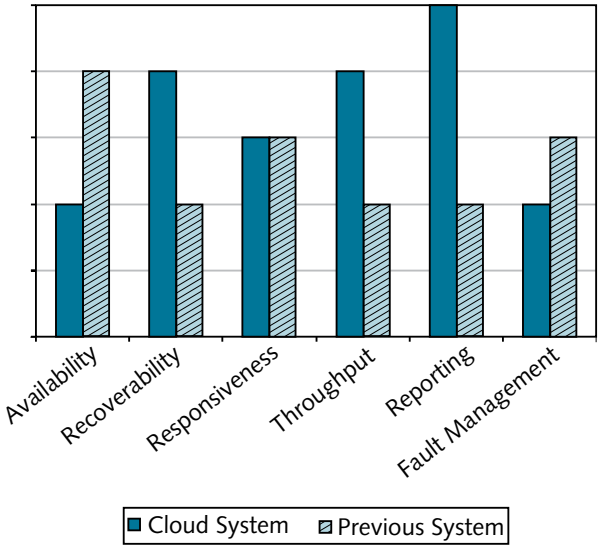

Figure 4.11: Performance and manageability comparison

The new cloud-based system is better in several areas. Recoverability and throughput have improved. The in-house system was struggling to cope with peak workload, and the availability of increased Capacity on Demand (CoD) has made a big difference to the performance of the system at busy times. And the reports from the service provider are very clear, and provide much better utilization information than the company had before. But there are two areas of concern. System downtime has been a problem, and the response of the supplier when faults have occurred has not been good.

Konsort-Prinz uses a similar approach to look at whether performance was within or outside contracted values and the SLA, as shown below (in Figure 4.12). (Good performance is indicated by points near the center, poor performance by points near the edge.)

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

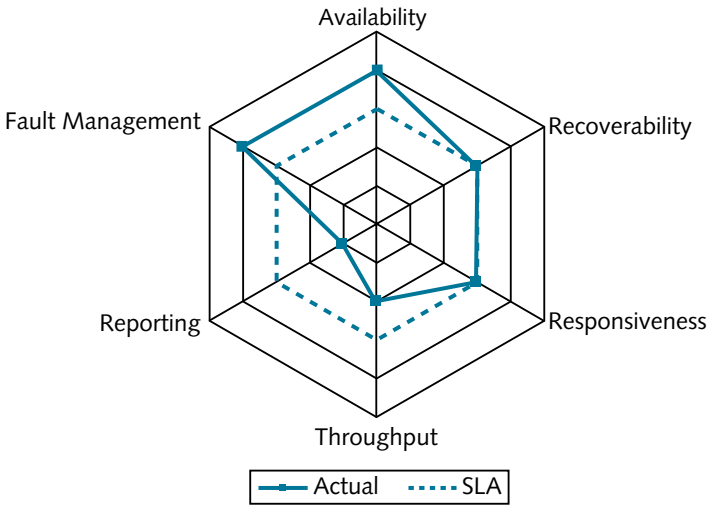

Figure 4.12: Comparison of actual performance with SLA

This confirms that supplier performance on availability and fault management has not been as good as promised. Konsort-Prinz will raise this with the supplier when the contract comes up for renewal, and will change supplier unless it receives satisfactory assurances that availability and fault management will improve.

ViWi's founder is not at this stage worried about performance and manageability. His main concerns as they move towards launch are with the financial model and with risk. These topics are for the next two chapters.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Chapter 5

Understanding cloud risk

The emergence of cloud computing is a fundamental shift towards new on-demand business models together with new implementation models for the applications portfolio, the infrastructure, and the data, as they are provisioned as virtual services using the cloud. These technological and commercial changes have an impact on current working practices. Businesses need to understand the impact of the new combinations of technology layers, and how they work together. A crucial part of this is analyzing and assessing the risks involved.

For example, the use of shared resources, in multi-tenanted cloud systems and across multiple organizations seeking economies of scale, results in companies relying upon a common cloud service or platform. What attendant risks might this bring to the tenant consumer of the service, and to the sellers and providers of the cloud services? How will it impact their expectations of service levels and performance?

This is a fundamental issue for any enterprise that considers using the cloud. As the Proposed Security Assessment and Authorization for US Government Cloud Computing points out: "The decision to embrace cloud computing technology is a risk-based decision, not a technology-based decision" (see [FEDRAMP]). All organizations, large and small, need to establish the right decision framework and governance mechanisms to use cloud computing successfully. These rely on an ability to analyze and assess the risks.

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). More information can be found on the CMU SEI Risk website [SEI RISK] and, in particular, in the excellent tutorial Rethinking Risk Management [D&A].

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

5.1 Risk management

Risk management is a core business activity of all enterprises, large and small. You can find a summary of the basic principles in, for example, the briefing document on Risk Management for SMEs produced by the Institute of Chartered Accountants in England and Wales [ICAEW]. The Committee of Sponsoring Organizations of the Treadway Commission [COSO] is the source of some of the important work on Enterprise Risk Management (ERM), which has developed into a significant discipline. This work, together with work done by national standards bodies (particularly in New Zealand and Australia) is the basis of ISO 31000:2009: Risk Management – Principles and Guidelines, which is now the standard risk management framework [ISO 31000].

Risk management is sometimes thought of as an unnecessary expense, a chore that can be neglected. This is the opposite pole of folly to the idea that risk can be avoided. A well-run business avoids the toxic risks, but accepts the manageable ones, takes care to minimize its exposure, and makes a healthy profit. Similarly, a well-run non-profit organization must know which risks to accept, and which to avoid, if it is to achieve its objectives. Good risk analysis often demonstrates that pre-existing worries may not actually turn out to be real risks. Properly applied, risk management creates and protects value.

The Mosaic approach builds on and extends traditional risk management to provide a framework for managing complex, systemic risks. It takes a holistic view of risk to objectives by examining the aggregate effects of multiple conditions and potential events. The issue of cloud risk is a collaborative and tenancy one. When an enterprise uses cloud computing, risks can arise in inter-related services that may not be under the complete control of the internal IT department. As [D&A] points out, traditional, tactical risk management is designed for environments with low uncertainty and few interconnections, but today's networked technologies operate in an environment of high uncertainty and dynamically changing, interconnected systems. This is particularly true for modern enterprises operating within cloud-based ecosystems.

5.1.1 Risk assessment

When a program to develop and deploy a new business solution is considered, you assess its *mission risks*: the systemic risks that affect the

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

program's ability to achieve its key objectives. A mission risk arises from a factor that has a strong influence on the eventual outcome or result. Such a factor is called a *driver*. Drivers enable a systemic approach to risk management by aggregating the effects of conditions and potential events.

You assess these risks initially, and reassess them periodically during solution development and operation. When the solution is in operation, the assessment incorporates measurements of solution parameters that affect the risk.

Assessing a mission risk typically means considering a number of complex, inter-related factors. The use of cloud computing affects many of these factors, sometimes negatively, sometimes in a positive way.

Assessment of mission risks is not a simple process. It may, for example, require interviews, documentation reviews, and group meetings. This clearly takes significant effort. Nevertheless, putting this effort into a systematic approach to risk management for a large and complex project will pay off. Even at the stage of preliminary assessment, where a large effort is not justified, it is worth following the general lines of the Mosaic approach when thinking through the risks involved.

This chapter discusses factors where the impact may be negative. There are many factors where cloud computing has a positive effect. For example, it can lower the risk that an enterprise cannot meet business demand due to lack of IT capacity. These advantages have already been covered in Why Cloud (Chapter 2).

5.1.2 Risk communication

Company executives are receiving requests for investment in projects with cloud computing and want to understand the ROI and investment risk. They seek to make value judgments and comparisons with available known performance benchmark data to assess the probability of success for a project. They want to understand any gaps or areas of risk that might need to be taken into account in decisions to approve a cloud project or in assessing the use of the cloud service.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

In a large enterprise, risk decisions are often taken jointly by several people, who then must explain them to the rest of the executive team. This implies a need to communicate risk information in a way that is clear and easy to understand. It may also be important to communicate risk information outside the organization; for example, to potential shareholders.

Risk management is an important technique used when implementing an enterprise architecture project (see Chapter 31: Risk Management of [TOGAF]). The enterprise architect will seek to satisfy the executives' needs for risk information with an outline analysis when setting the architecture vision, and a detailed analysis when developing the business architecture.

You may be conducting the risk analysis, and be asked to explain it to others. More likely, the detailed analysis will be done by others, and you will want them to explain it to you. You need to understand how to communicate risk information, either to do it yourself, or to insist that others do it in a way that you can easily understand. The approach followed here, based on Mosaic, uses simple scorecards that communicate the risks very clearly.

The *impact* and *probability* of each risk are assessed separately, and then combined to give an indication of *exposure*. A scale with five levels is used for each of these quantities:

1. Minimal
2. Low
3. Moderate
4. High
5. Maximum (or Critical)

This enables a clear and direct graphical representation of project risks. For example, the risk exposure summary in the figure below (Figure 5.1) is taken from the initial risk analysis on the ViWi project. (This is one particular example. The impact, probability, and exposure values will differ in other cases.)

It shows major exposures to financial, disaster recovery, and system quality risks. The financial and system quality risks are inherent in the project because no-one can be sure that people will buy the virtual widgets or find them useful. The disaster recovery risk can, however, be mitigated by data

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

back-up arrangements. The analysis makes it clear to the decision-makers what risks they are taking, which of them can be mitigated, and which must be accepted. That is its value.

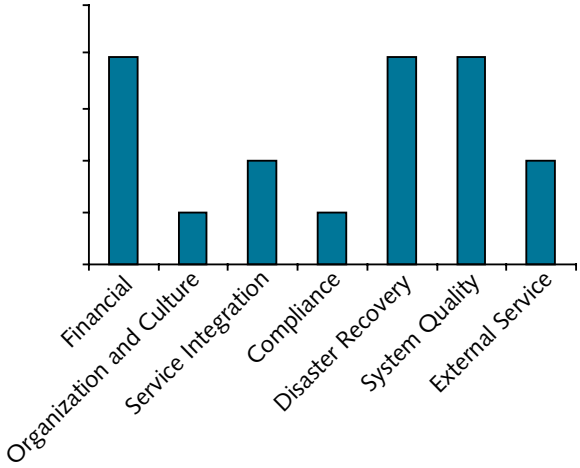

Figure 5.1: ViWi risk exposure

5.2 Cloud mission risks

The main cloud-related mission risks to consider are the ones shown in the ViWi example above:

- The solution does not meet its financial objectives.
- The solution does not work in the context of the user enterprise's organization and culture.
- The solution cannot be developed due to the difficulty of integrating the cloud services involved.
- The solution does not comply with its legal, contractual, and moral obligations.
- A disaster occurs from which the solution cannot recover.
- An external cloud service used by the solution is inadequate.
- The system quality of the solution is inadequate, so that it does not meet its users' needs.

It is important to have a clear definition of each risk, and also to explain the factors involved and provide rationale for each assessment. The figure below (Figure 5.2) shows an outline of the rationale for Sam Pan Engineering's initial assessment. A full assessment should have more detail. While the

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

scorecard chart communicates an overview of the risk clearly, people will often need more information before taking specific action. And it is crucial that repeated risk assessments are conducted on the same basis. This is impossible without clear explanations of the factors and decision rationale.

Factor	Impact	Impact Rationale	Probability	Probablility Rationale	Exposure
Financial	4	Although there will be serious consequences if the system does not meet its financial objectives, it may still play a useful role within the corporation	3	The ROI calculations make a number of assumptions, some of which are quite likely to be unfounded.	2
Organization and Culture	5	Critical mission risk	3	People are ready and able to cope with change The constituent companies of the group are suspicious of central control and like to be independent These companies have some budgetary discretion and can buy their own IT Group IT cannot instruct constituent companies to adopt a particular solution	3
Compliance	5	Critical mission risk	1	The group does not operate in countries that currently have legislation that would affect the system. There is a possibility of such legislation being brought in at some point in the future.	1
Disaster Recovery	3	Distributed system supporting many projects over a wide geographical area - a disaster might affect some of these, but not a major proportion	4	A physical disaster that affects part of the system is quite likely to happen sometime somewhere. (Some company offices were damaged by the 2004 tsunami, for example.)	2
System Quality	5	Critical mission risk	3	See separate assessment	3

Figure 5.2: Sam Pan Engineering initial risk assessment

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Note that Sam Pan Engineering has not assessed the *service integration* and *external service* risks, because it is not using any external cloud services.

Note also that system quality is a complex risk that is the subject of its own separate probability assessment. The result of this is carried forward into the overall assessment.

5.2.1 Financial

That the solution will not meet its financial objectives is a mission risk. Many promising solutions have been abandoned because they fail to meet their revenue and profit objectives. The impact of financial risk is almost always critical.

Business models generally look great on paper, but the figures can be very different when they are put to the test. By enabling costs to be related directly to workload, and hence to revenue, cloud computing reduces some of the risk. But it affects other factors differently, sometimes in a negative way.

The key factors to consider when assessing cloud ROI risk probability are the leading indicators discussed in the next chapter: utilization, speed, scale, and quality. These factors are built into most ROI models, and affect the headline figures for investment, revenue, cost, and time to return. Deviations of these figures from what was planned indicate likely deviations of the headline figures, and consequent failure to achieve the desired ROI.

The impact of mission risks is usually critical. Sam Pan Engineering, however, assesses the financial risk impact only as being high. Although there will be serious consequences if the system does not meet its financial objectives, it may still play a useful role within the corporation.

5.2.2 Organization and culture

Introducing cloud computing may require major changes to an enterprise's organization and culture. Can you persuade a traditionally run company with a 'we do it our way' culture to use external, standardized services? Will you have to make the entire IT department redundant? What will the unions do?

A company adopting new cloud-enabled business processes must have a clear executive vision and direction for business transformation. The program

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

must have the support of top-level executives, and they must share a clear vision of what is to be achieved, and define a clear strategy for realizing it. This should include:

- Establishing a clear roadmap for procurement or implementation of cloud services and applications that use them
- Coordination of stakeholders and competing strategies to get consensus for storage, computing, network, and applications to avoid islands of demand usage
- Identification of pilots and demonstrations to create confidence and build buy-in and usage in the user community for cloud services
- Building transition skills and cloud technology knowledge in-house or with service partners to ensure the technical risks are understood and managed

The organization must be able to make the change to adopt on-demand provisioning fast enough to gain the consequent business benefits. Cloud aims to be a step change in total speed of delivery. The organization must be able to move up a gear to take advantage of this.

The organization's financial structure, and in particular its charging methods, may need to be redefined to adapt to cloud service business models.

Traditional working practices may need to change. This does not just affect the IT department. The fact that the infrastructure or applications used by business departments is not directly controlled by the organization may mean changes to the way that those departments work. This will require education and training, and may meet with resistance from the workforce.

All of our example companies rate the impact of this risk as critical. ViWi rates the probability as minimal; as a start-up company created for this project, it has no organizational or cultural barriers. The other two companies believe they have moderate organization and cultural risk probabilities.

5.2.3 Service integration

As use of cloud computing develops, it is increasingly likely that an enterprise will use not one but several cloud services, and will need to integrate them with each other and with in-house systems. For example, in Case 21: Brand Unification in Cloud Computing in Use (Appendix A) the enterprise wants to

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

use cloud-based collaboration services together with other services replacing some of its existing applications. These services must be integrated with each other and with those of the existing applications that are not replaced.

There is a risk that it will not be possible to integrate the cloud services with the existing system and with each other. This risk is critical; if the system cannot be built, it cannot be used. The service integration risk can be assessed by considering interface conversion cost, ability to change the existing system, and available skills.

The interface conversion cost is the cost of providing 'glue' software to connect the services if their interfaces are not completely compatible. The interfaces may be compatible if the services are designed as part of a common suite, or to fit within a particular industry reference model. Otherwise, you should look at how far they have compatible syntax, compatible semantics, and fit within a common process model, or support a suitable interoperability protocol and architecture.

Usually, syntactic conversion is relatively straightforward, and semantic conversion is possible but expensive, but if the services have radically different process models then integration may be extremely difficult.

Similar considerations apply to the ability to change the existing system. You should assess whether it can be used with little change, or whether radical redevelopment is required. In the latter case, the risks are high.

Significant skills are required to assemble and customize multiple cloud services from different providers in a flexible, adaptable way, while maintaining security, back-up, and governance mechanisms. This assembly of services will mean that applications using them will need to become more 'loosely coupled' – programmed to act with an integration layer, not the underlying infrastructure. You should assess how far you have these skills in-house, and the cost of employing specialist consultants. If major changes are needed to legacy systems, you may find that the necessary skills no longer exist, either inside or outside your organization.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

5.2.4 Compliance

The requirement to meet compliance obligations was discussed under Establishing Requirements (in Section 4.2). For each of your obligations, you should assess the risk that the obligation will not be met.

In addition to assessing risks associated with your own systems, you should assess risks associated with external services that you use. For projects based on cloud services, this is often the most significant source of compliance risk.

Regulations, or company policy, may require data to be located in particular geographical areas or legal jurisdictions, to be kept secure, with its integrity and confidentiality maintained, or to be kept online or archived for specific periods. This applies particularly to personal and financial data.

The impact of failing to conform to these regulations varies considerably. For legal requirements, it depends on the prescribed penalties and the enforcement regime. For moral obligations, the impact can include loss of reputation and standing, which may in turn be reflected in market share.

Dependence on an external cloud supplier can increase the probability of non-compliance. Even if you have contracts that provide the necessary assurances on location and confidentiality, *force majeure* may prevent the supplier from honouring them. For example, what would be the result of legal action for subpoena of data in a cloud environment that may not even be held under your tenancy, but has been placed on the same system by other tenants? And what would then be the impact on your corporate reputation?

Because of questions like this, compliance is an important cloud risk area. All of our example companies rate the impact as critical, but Konsort-Prinz is the only one that gives it a significant probability. The company has to conform to European legislation. It believes that it can do this, but it is a complex area, and the company rates the risk probability as moderate.

5.2.5 Business continuity management

It is sometimes necessary to react to and recover from unplanned events. As for compliance, you should assess risks associated with external services that you use, as well as assessing risks associated with your own systems.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Physical disasters such as fires, floods, and hurricanes that result in a major loss of IT capability are examples of events that can affect in-house systems or the systems of cloud suppliers. And there are other kinds of event, such as mergers and acquisitions of suppliers, unforeseen bankruptcy, or cancellation of contract, that can affect a supplier and require a similar response.

Using cloud computing can make it harder to respond, because you have less control or less visibility over the IT resources concerned.

As part of your risk analysis, you should identify the unplanned events that could harm you, and assess their probabilities and impacts. You may also wish to make general provision for unforeseen events that disrupt the cloud services that you use, or damage their data.

Having identified the risks, you can build into your system design elements that will reduce their probability or mitigate their effects. For example, an effective back-up and restore process, with the back-up copy held in a different location from the data, or on your own rather than the cloud supplier's system, can change the impact of a disaster from fatal to merely serious.

5.2.6 System quality

This is the risk that the solution does not meet its users' needs. The importance of system quality is discussed under Quality: Improved Margin from Better Service (in Section 6.4). The impact of a shortfall is reduced margin and loss of ROI.

As for compliance and disaster recovery risks, you should assess system quality risks associated with external services that you use, as well as assessing the risk associated with your own systems.

The risk factors for system quality are illustrated in the first system quality assessment carried out by ViWi. Unsurprisingly, since the product has not yet been tried in the market, this shows high failure probabilities for the high-impact factors of functionality and user satisfaction. (Distance from the center of the chart indicates the size of the impact or probability.) If these risks materialize, and the product does not satisfy its users, then the project will fail, and its backers will lose their money.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

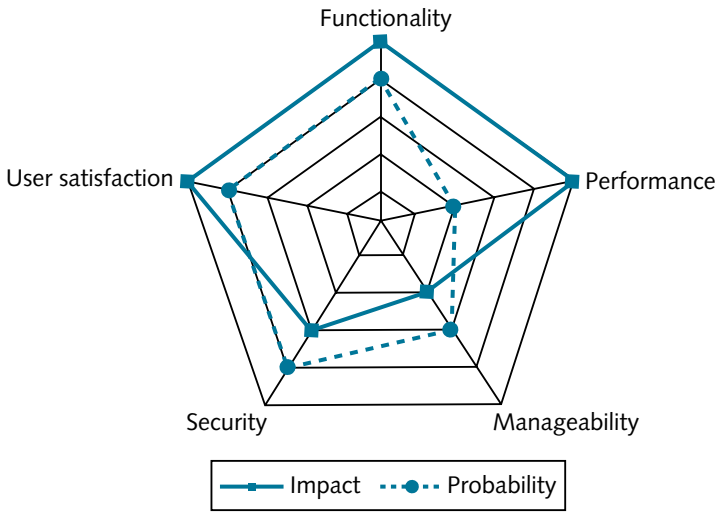

Figure 5.3: ViWi initial system quality assessment

System quality risk factors are discussed in more detail in a separate section (Section 5.3).

5.2.7 External service

If you are using external cloud services, then you should assess the risks that they will not be adequate.

The impact can be critical if your solution depends on an external cloud service for essential infrastructure or functionality. But the impact is less if you have an exit strategy, as discussed under Selection (in Section 4.3).

The factors to consider include the system quality of the external service, and the adequacy of its supplier.

The system quality of an external service can be assessed using the same factors as for the system quality of your own solution. The assessment made by ViWi of the risks associated with using an external PaaS service is illustrated below (in Figure 5.4).

Although the factors are the same as for the assessment of the overall ViWi solution, the values of those factors are different. They are often quite unrelated. For example, the users of the PaaS service are the ViWi developers, and it is not so critical for them to be satisfied with the development platform

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

as it is for the virtual widget users to be satisfied with the end product. Also, the PaaS platform is well established, and the probability that it will be unsatisfactory is low.

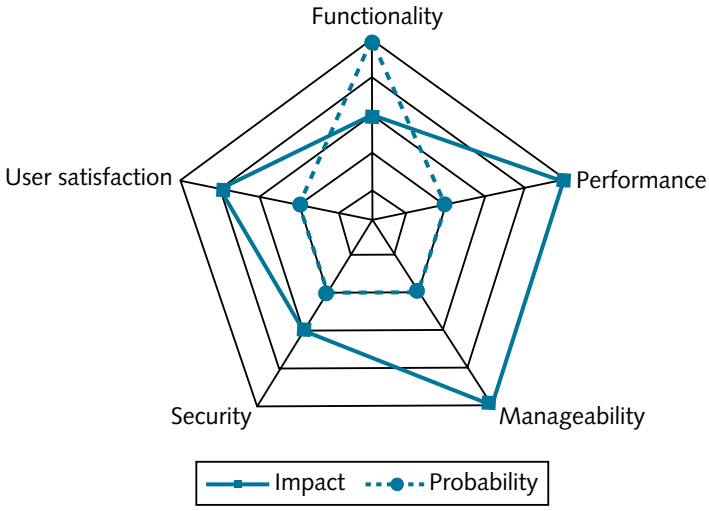

Figure 5.4: ViWi external cloud service system quality assessment

Enterprises that use cloud computing rely on cloud suppliers for many things that traditionally would be done in-house. You should assess supplier quality because there is a risk that a cloud supplier is not adequate; for example, that it proves untrustworthy, or goes bankrupt.

Many of the factors involved – such as financial stability, size, reputation, and track record – are similar to those for any other kind of supplier on which you rely. In the context of cloud computing, you should look particularly at the supplier's track record in meeting SLAs, responding to complaints, and being willing to share information about service operation and system architecture.

Remember, also, that 'possession is nine points of the law' when it comes to information. Providers can, for example, use or sell subscriber information for marketing purposes in ways that can be hard to detect or prove. Supplier quality assessments should include trustworthiness.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

5.3 System quality risk factors

System quality as defined above depends on the factors of *functionality*, *performance*, *manageability*, *security*, and *user satisfaction*. The requirements for the first four of these factors are discussed under Establishing Requirements (in Section 4.2), and user satisfaction depends on those four factors.

5.3.1 Functionality

The risk that the system does not have the necessary functionality depends on the similar risk for external cloud-based systems that you use, as well as on non-cloud-related factors, such as the quality of the specification.

If you are relying on PaaS or SaaS services with complex functionality, this may be a significant consideration.

ViWi is relying on significant functionality in the cloud platform that it uses. The company has already found that there is no PaaS service that meets its requirements, so the risk probability is certain. The impact is, however, only moderate; it can overcome the lack of platform functionality by writing additional software in its own system.

5.3.2 Performance

Performance covers *availability and reliability, recoverability, throughput,* and *responsiveness*.

Availability depends on reliability. From a risk analysis perspective, the MTBF reliability factor determines the risk probability, and the MTTR – the probable value of the time needed for repair of the system – determines its impact.

Fault tolerance is a factor that affects the availability risk probability. A fully fault-tolerant design has no Single Point of Failure (SPOF), and often accommodates multiple failures within a service window. Software fault-tolerant designs include exception handling and task rollback. For such systems, the failure probability will be low.

Recoverability is the ability to recover from failure. If a system fails, there is likely to be some data loss. If the failure is to a processing unit, the amount

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

of data lost will depend on how the program handles the data; with well designed software, the amount can be small. If the failure is in a storage system, the loss may be significant. Use of redundant storage can minimize the probability of this.

Many companies would go out of business if they lost all of the data that they keep on IT systems. It is normal practice to limit this impact by taking back-ups. For example, a daily back-up limits the impact to loss of 24 hours' data.

If the system is performing adequately, its throughput is the same as the offered workload. But there is a risk that it cannot handle the load required. In assessing this risk, you should consider the range of throughput levels that the system should support, the degree of predictability of the variations in level, and how changes in load will be handled.

If you are using cloud services, you can make a trade-off between risk and cost. Configuring your system with more resources than are needed to meet the expected load reduces the probability that the load cannot be handled, but increases the cost. If you can configure the system dynamically, you can keep this over-provisioning to a minimum, but must consider how quickly the load might change, and how quickly additional resources can be brought online.

The probability of an overload can be reduced by apportioning the offered load between systems. Cloud providers normally have some means of doing this when resources are provisioned, and Sam Pan Engineering will have this facility in its private cloud platform. Konsort-Prinz uses a single central database, and cannot partition the processing. ViWi can easily do so, because different virtual widgets can run on different systems.

In assessing the risk that a system is not sufficiently responsive, consider the response times that it should have, the permissible variability in response time, and the need for predictability of response.

Poor response time is often correlated with throughput overload. We will see an example of this in the next chapter.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

5.3.3 Manageability

Manageability includes the factors of configurability, reporting, and fault management. The requirements for them were discussed under Establishing Requirements (in Section 4.2). Failure to meet these requirements leads to user frustration and delay at best, and at worst renders a system completely unusable.

Provisioning management is particularly important for cloud services. Konsort-Prinz assesses the manageability risk impact of the cloud service in its solution as high, principally because of the needs of its twice-daily provisioning exercise. ViWi rates manageability of the platform as critical, because of its need for dynamic resource provisioning.

5.3.4 Security

Having your own information, on your own hardware, and between your own four walls, provides a level of comfort that you lose in the cloud. Cloud services are often accessed over the public Internet and this must be considered when assessing security. This is not to say that cloud computing is necessarily insecure, just that new considerations need to be taken into account and more modern security models developed and applied. You must adapt traditional security models to suit cloud computing needs and consider end-to-end security, including your own internal policies for access control and user provisioning.

The analysis of risks associated with security threats is an important and specialized discipline. This Guide does not attempt such an analysis. It provides general advice on assessing cloud security risks for the purpose of business decision-making, but contains no information on identifying or taking countermeasures against specific security threats. You can find detailed information on risk in the context of data security in The Open Group Risk Taxonomy [RISK] and The Open Group FAIR – ISO/IEC 27005 Cookbook [FAIR]. The Cloud Security Alliance [CSA] has produced some excellent material relating to security in the context of cloud computing.

The requirements for end-user access control, supplier access control, resource partitioning, security logging, and threat response were discussed under Establishing Requirements (in Section 4.2). Failure to meet these requirements can lead to system and data unavailability, financial loss, leakage

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

of sensitive information, failure to meet privacy regulations, and damage to reputation. It can also result in the integrity of the system being compromised, leading to sabotage or data leakage. This can be achieved by insertion of malicious software or by abusing existing software vulnerabilities, for instance.

There are two factors that can increase the probability of a security breach for consumers of cloud services. One is the need to rely on the service provider for part of your security arrangements. The other is that the security arrangements in a system that uses services from different providers are likely to be complex, and this introduces the possibility of unconsidered areas where there can be gaps in the defences.

Konsort-Prinz assesses the impact of a security breach as high, in view of the possible disruption to its operations and loss of customer information, but not critical. It assesses the probability as low, because it trusts the supplier, only its staff can control operations and access customer information, and it will have standard access control mechanisms in place. If it decides to give its suppliers access to the system, to boost ecosystem productivity, the security risk must be reassessed.

5.3.5 User satisfaction

Logically, user satisfaction is a consequence of the other system qualities considered above. Practically, it should be considered and assessed separately. There is a definite risk that the users are not satisfied with a system, even though it appears to have adequate functionality, performance, manageability, and security.

At design time, assess the probability of user satisfaction by looking at whether human factor experts have been consulted, and whether user trials have been done. When the system is in service, take proactive steps to measure user satisfaction. The importance of this as an indicator of quality is discussed under Measuring and Tracking ROI (in Section 6.6).

5.4 Continuing risk assessment

An initial risk assessment is made when taking a decision to use cloud computing, or deciding which form of cloud computing to use. The examples given so far in this chapter have been of such assessments. But the

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

risk analysis is not just done once at the start of system development and then forgotten. It should be repeated throughout the life of the system as circumstances change.

5.4.1 Solution architecture development

The importance of risk management as an integral part of enterprise architecture is described in Chapter 31: Risk Management of [TOGAF]. This section illustrates the application of this principle to the Konsort-Prinz cloud solution.

The initial risk assessment for this project indicated significant exposure to risks in the areas of organization and culture, compliance, and disaster recovery.

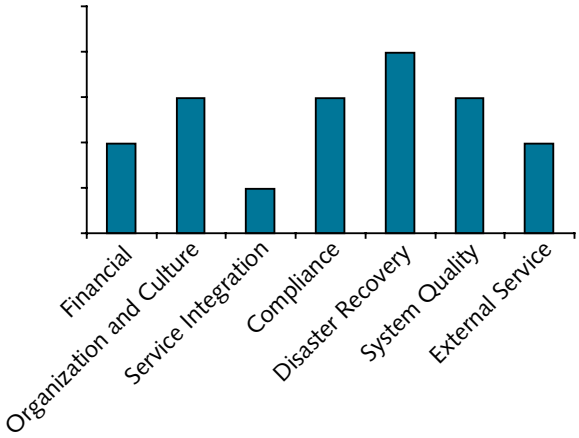

Figure 5.5: Konsort-Prinz initial risk assessment

The organization and culture risks arise partly because of a generally conservative culture within the company, and partly because national legal constraints make it difficult to make people redundant without their consent. However, the project has the backing of top-level management. The CEO will make it plain that opposition will not be tolerated. At the same time, the company will offer financial terms that will encourage the IT staff who are no longer needed to leave.

The compliance risks relate to European data protection legislation. This is complex and still evolving. (There is a central EU Directive, but it is its interpretation in each member state that counts.) Konsort-Prinz decides to

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

keep all data within the EU, to ensure that all customer and employee data is kept confidential, and to enable customers and employees to check their personal data. The executives believe that this is a moral obligation, and also that it will mean that the company complies with all applicable laws. The implications for the IaaS service that the solution will use are that it must keep all data within the EU, and make that data accessible only to people authorized by Konsort-Prinz. These are made mandatory requirements in the selection process.

To provide for disaster recovery, there will be a back-up IaaS provider. Copies of all the data will be stored periodically on this provider's system, and there will be a contingency plan to transfer operation to that provider if the main provider suffers a disaster. The impact of a disaster will then be at a non-critical level, although there may still be disruption to operations and some data loss.

The system quality risk is in the performance area. A failure just before the daily back-up during the seasonal peak could cause a data loss that would affect 5% of the year's sales – which would be reflected quite visibly in the bottom line. The designers decide to modify the system to write a copy of the transaction data to a separate storage resource within the cloud provider's system. This is a good example of how an effective risk assessment can find a small 'weakest link' within a large, and seemingly strong, chain.

The effect of this risk mitigation process is to reduce all risk exposures to an acceptable level.

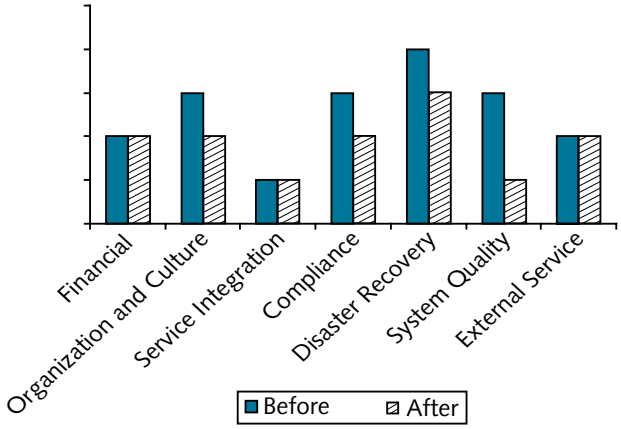

Figure 5.6: Konsort-Prinz risk mitigation

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

The assessment will be repeated at significant stages of the architecture development to ensure that the levels of risk exposure continue to be acceptable.

5.4.2 Cloud service selection and procurement

When cloud services are selected, the risks associated with each should be reviewed. Negotiation of contract terms should seek to reduce the risk to the purchaser. The level of risk remaining after this negotiation is a consideration to take into account when making the choice.

Konsort-Prinz has identified three IaaS suppliers that meet the requirements and whose prices give reasonable cost models. (The prices were taken from the public websites of three real IaaS providers in October 2010.)

	IaaS 1	IaaS 2	IaaS 3
Price per processing unit per month	$68	$43.20	$55.00
Price per storage Gigabyte per month	$0.20	$0.25	$0.15
Price per Gigabyte IO	$0.13	$0.00	$0.20
Overprovisioning factor	2.5	4.0	2.5
Processing cost per year	$427,500	$432,000	$343,750
Storage cost per year	$62,862	$78,578	$47,147
IO cost per year	$97,500	$0	$150,000
Annual cost	$587,862	$510,578	$540,897

Figure 5.7: Konsort-Prinz supplier cost models

The first is a large multi-national corporation. The second is a small local company that offers extremely cheap rates on monthly contracts (hence the need for a larger over-provisioning factor). The third is a medium-sized company based in another EU country.

Konsort-Prinz makes a risk assessment for each potential supplier, and compares the results. The system quality risks are pretty even, but the risk that the suppliers are inadequate differ: minimal for the large company, moderate for the small local company, and low for the medium-sized EU company. In spite of its cost advantage, Konsort-Prinz assesses the risk for the local supplier (IaaS 2) as being too great. The larger EU company (IaaS 3) is slightly riskier than the multi-national, but is selected on cost grounds.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

5.4.3 Solution operation

Risks should also be reassessed as appropriate during operation. Of our example companies, the best example of the value of this comes from ViWi.

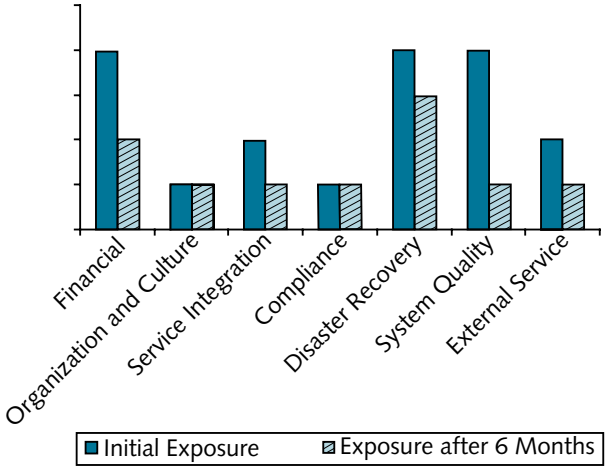

Figure 5.8: ViWi six-month risk assessment

The risks are reassessed six months into the project. The product has been launched, and the question marks over service integration, system quality, and the external service have been resolved. In addition, although sales are lower than expected, the ROI factors have improved. (This will be explained in the next chapter.) The overall risk picture is very different, and much more positive.

In this case the assessment is crucial for a decision on whether to continue the project. More usually, the stakes are lower. Nevertheless, risks do change over project lifetimes. This may mean system modifications and changes in operation, even though cancellation is not in question.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Chapter 6

Building ROI from cloud computing

Many businesses have addressed emerging trends through actions to drive cost reduction and leverage IT service providers' adoption of cloud-style services. Many industry organizations and leading IT suppliers of software, hardware, and services, seeking to address their customer needs, have vigorously evaluated and followed a cloud-style strategy. The challenges and issues are in the transition from the current traditional IT to the new potential capabilities of cloud computing. They must be expressed in a language that end users can understand, and relate to investment, cost improvements, or business performance.

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. How do you do this? What are the qualities of cloud computing that affect ROI?

ROI is the proportionate increase in the value of an investment over a period of time. It can be measured in a variety of different ways, but there are just four basic ways to improve it: decrease the investment, increase revenue, decrease costs, and make the return faster. Using cloud computing, you can achieve any of these; but you cannot achieve them all at the same time.

It is the relationship between the factors that counts, rather than the absolute values. If you move to public cloud, you generally decrease investment but increase cost. With private cloud, it is the other way round. You can improve or worsen ROI either way, depending on your revenue and speed of return. Revenue can be increased by improving the delivered features and quality – which enable a higher price to be charged – or by operating on a larger scale. But improvements in features, quality, and scale generally also mean higher costs. You must get the balance right.

Just looking at cloud computing from a technical infrastructure point of view is potentially missing the wider picture of the impact of technology on the

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

business. Overall, what matters is defining the value to business. Value can be defined in many ways. It does not just mean the financial values, but can also mean customer value, seller value, broker value, market brand value, and corporate value. These aspects should not be neglected.

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 relate to *productivity*, *speed*, *size*, and *quality*. This chapter describes how each of these drivers contributes to ROI, and shows how to use them to compare cloud and traditional IT solutions, and how to monitor them to maintain and build ROI from cloud computing.

6.1 Productivity: more business with less IT

The capacity/utilization curve used by Amazon Web Services to illustrate the advantages of cloud computing is often quoted. The figure below (Figure 6.1) uses a similar curve to illustrate the central idea of utility-based services enabled through on-demand provisioning to meet actual usage. The model illustrates the central idea around utility-based services enabled through on-demand provisioning to meet actual usage.

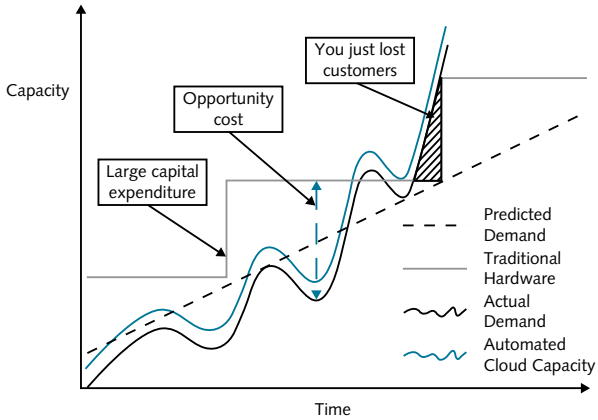

Based on illustrations by Amazon Web Services

Figure 6.1: The capacity *versus* utilization curve

The effect is equally clear in the capacity-utilization chart for Konsort-Prinz's first year of use of its new cloud-based system, shown earlier (in Figure 4.8).

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

For most of the year, the in-house system would have been used at a small fraction of its capacity, but in the peak month it would not have had sufficient capacity, and business would have been lost.

6.1.1 Resource utilization

Resources that are sized to handle peak loads are under-utilized at off-peak times. Also, in many enterprises, servers are dedicated to specific functions or departments, and can be massively under-utilized. It is commonplace for departments to request new servers for new projects, even though other departments have capacity that is going begging.

Cloud computing enables resources to be shared by different loads, and thus improves utilization. The sharing can be between enterprises, with public or community cloud, or within an enterprise, with private cloud. The rewards can be high. As an indication of what can be achieved, enterprises that adopt private cloud as a solution are reporting server consolidation ratios of as much as 12 to 1.

At the infrastructure level, resource sharing means supporting many platforms and applications on the same physical resources, using the virtualization techniques described under IaaS in the section on Service Models (Section 1.3).

At the platform and application levels, providers can achieve resource sharing through multi-tenancy. For PaaS, this means applications from many clients running on the same operating system. For SaaS, it means clients sharing the same application instance. By doing this, using specially-designed application architectures, they can achieve better utilization of the underlying assets than if each client uses either a separate application instance on the same operating system or another virtualized instance sharing hardware.

6.1.2 Usage-based pricing

Usage-based pricing translates the higher utilization achieved by providers into lower costs for consumers.

The benefit to Konsort-Prinz is not that it is using less resource, it is that it is *paying for* less resource. Its cost curve follows its resource utilization, resulting in a lower total than for a dedicated, under-used system.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Usage-based pricing for software licences gives similar benefits to SaaS users. Traditional licensing associated with ownership, number of users, support, and maintenance costs and services leads to capacity-utilization gaps similar to those for hardware resources.

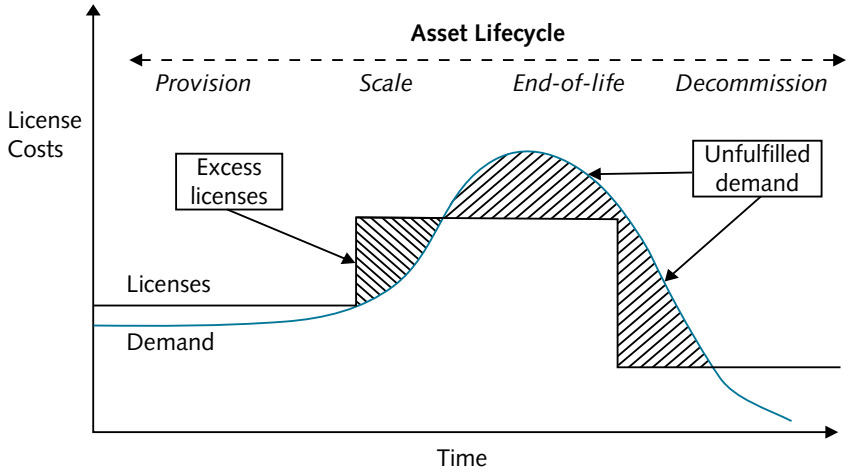

Figure 6.2: Software license costs

These gaps are being addressed by the pay-as-you-go principle of on-demand cloud computing, which is changing the software cost model.

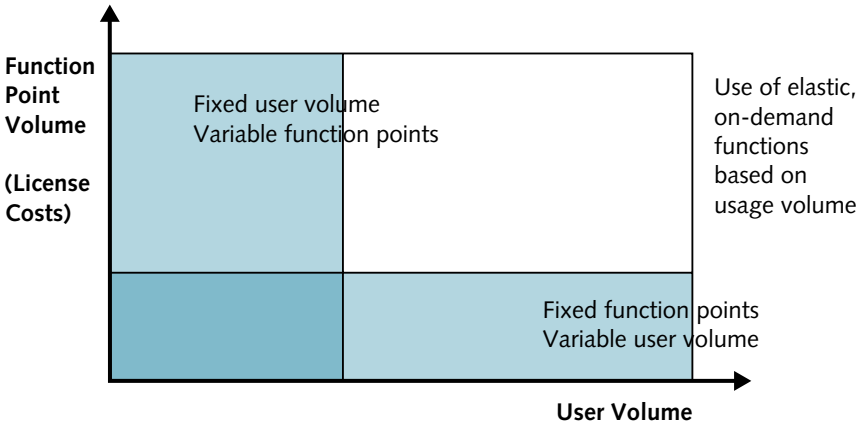

Figure 6.3: Software cost optimization

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

6.1.3 Specialization and scale

In addition to the advantages of load sharing, cloud computing can result in lower IT costs because of skill specialization and economies of scale. A large cloud provider – or private cloud division within a corporation – can be much better at providing IT services than a small IT department. And it can amortize the cost of problem-solving over a larger user base: the problems experienced by one user can be proactively fixed for all the other users of the cloud service, once they have been solved.

6.2 Speed: getting there more quickly

The second major way in which cloud computing contributes to ROI is through improved speed of operation.

6.2.1 Time to deployment

Cloud computing provides an increase in provisioning speed, which enables enterprises to acquire the resources they need faster. And, because resource configurations are visible, it speeds up the choice of multi-sourced resources, which can dramatically cut the time to deployment of new products and services. This applies to application and platform resources, as well as to physical infrastructure.

Elastic provisioning creates a new way for enterprises to scale their IT to enable business to expand. The compression of provisioning time from weeks to hours, for example, demonstrated by cloud computing providers is a means to rapid execution that is not just about saving time, but is also defining a new business operating model.

Organizations can review and develop business plans and then deploy infrastructure and services in a more rapid and proactive way. Customization and development, testing, and support can also be seen in a new light with the provision of IT services in a dynamic fashion targeted at business needs.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

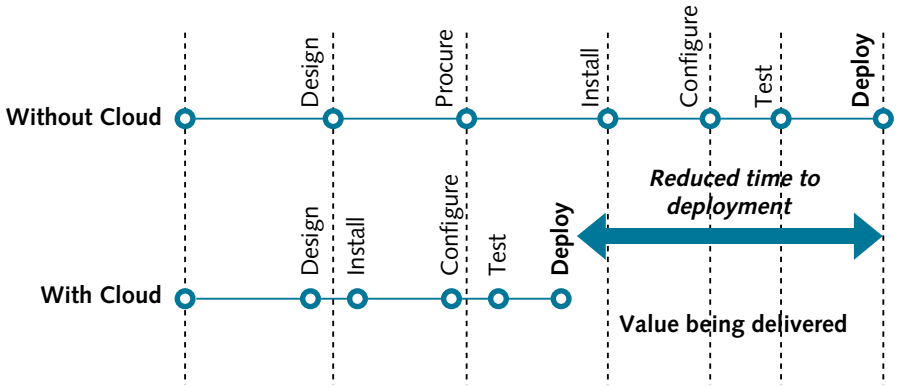

Figure 6.4: Time to deployment

ViWi provides a prime example of the ability of cloud computing to cut the time-to-market. The company's founder was a college whizz-kid with a great idea. Using cloud services – which were free for the low usage that was all he wanted – he could quickly build a proof-of-concept program to demonstrate the idea to potential backers. With a ready-to-use cloud development platform, there was no time wasted on buying, installing, and setting up hardware and software. Production versions could be tested in their deployment environment on the cloud, with no need to create special test-bed systems. The time from concept to launch was dramatically reduced, enabling the product to keep ahead of potential competition.

6.2.2 Lifetime cost models

Increased speed of execution has a positive impact on lifetime cost models. Typically, cost is reduced over the lifetime of a product or service as the depreciation cost of purchased assets decreases and as efficiencies are introduced. The speed of cost reduction can be much higher using cloud computing than traditional investment and divestment of IT assets, as illustrated in the figure below (Figure 6.5).

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

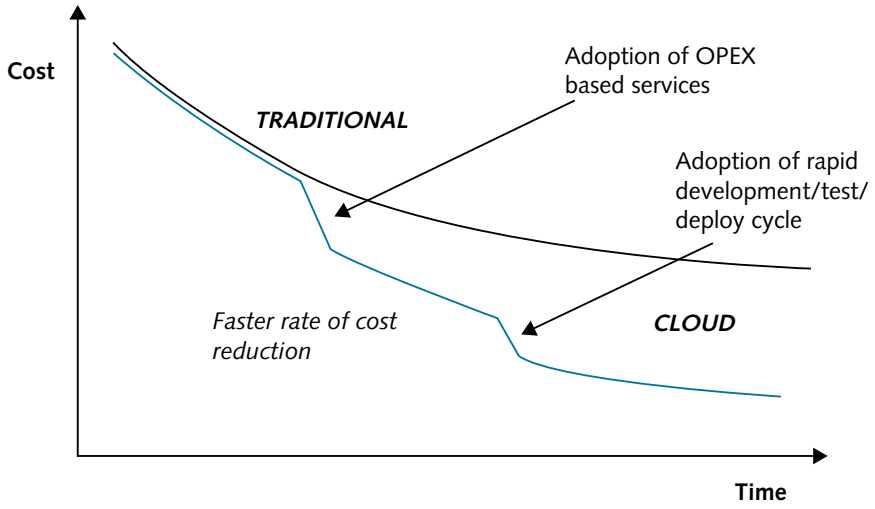

Figure 6.5: Speed of cost reduction

This higher rate of cost reduction means that profitability increases more quickly, giving shorter pay-back times and increased ROI.

6.2.3 IT asset management

As well as accelerating business process execution, cloud computing speeds up the process of IT asset management.

The use of IT has become an enduring feature in all organizations today. The investment in data, knowledge, and infrastructure assets and software now represent lifeblood operations for many businesses. But many enterprises have IT asset portfolios that do not reflect their real requirements. The issues of cost of ownership are often decoupled from choices made during selection of new IT, and the impact on the long-term running and maintenance of these IT services and subsequent business usage is not properly considered.

Technology design choices and purchasing are often done by strategic or tactical contractual purchasing based on project requirements, with little consideration for optimizing running and maintenance over the whole system lifecycle, but the costs of maintenance and modifications often represent a significant part of the asset lifecycle beyond initial provisioning.

The ability to 'design and provision for run', so that the choices of IT procurement are aligned with the best options and performance for

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

long-term operation, has long been an ideal goal of business and IT. But, while technical trends such as object-orientation, SOA, and Web 2.0 have brought functional improvements, the improvements in runtime infrastructure support have always been elusive.

A key aspect of moving to cloud computing is the ability to select hardware, software, and services from defined design configurations to run in production. It bridges the design-time and run-time divide and optimizes service performance. Patches and upgrades to the cloud resources are provided as part of the service.

Cloud computing facilitates optimization of the total asset portfolio. It can help an enterprise achieve the goal of a more cost-effective asset–management lifecycle process for the IT portfolio, to optimize both design and run-time performance.

6.3 Size: breaking new ground

The combination of lower cost and faster delivery makes products and services more competitive, which generates more business and leads to a larger scale of operation. But what makes cloud computing really exciting is that its potential effect on business is not just incremental improvement but disruptive transformation, which it enables through new operating models.

6.3.1 Entering new markets

Cloud computing removes the need for additional infrastructure to test and enter markets (a key advantage particularly for small to medium-sized organizations). It enables businesses to pursue new and existing markets by rapid entry and exit of the products and services. It enables enterprises to 'land and expand' in markets with an infrastructure and service capacity that can grow with the business. Existing and new markets can be attacked and entered through speculative and well-timed interventions to exploit and grow business performance.

6.3.2 High-value services

Internet services such as Facebook, MySpace, Flickr, YouTube, and Twitter have become part of our everyday social interaction, and businesses are increasingly making use of IT-based collaboration services for

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

communication, information exchange, and virtual meetings. In many cases, these are not provided by in-house IT; they are cloud services. Use of these services can transform business operation.

From a provider's point of view, these are prime examples of simple, high-value services that are massively successful. But they do not have traditional price-based revenue models; they are often free to use. Advertising is one way in which such services can obtain revenue. Another is through offering additional 'annuity' services – essential or useful add-ons that are paid for; just as many printer companies make money from the ink cartridges rather than from the printers themselves.

6.3.3 The long tail

At the other end of the scope and complexity scale, there are many cases of companies large and small that can enter and develop service offerings through a *long tail* approach enabled by cloud computing. For any product, most people will be satisfied by a small set of standard features, but there is a 'long tail' of potential users that require customization. The ready availability of machine intelligence in the cloud enables this customization to be delivered at a lower cost, so that serving the long tail becomes an economic proposition. (See The Long Tail by Chris Anderson [LONG TAIL].)

6.3.4 Becoming a cloud provider

The ability to become a cloud provider is one way in which new business opportunities arise.

Case 16: Service Delivery in Cloud Computing in Use (Appendix A) provides an interesting example of this; the telecommunications service provider will provide 'network as a service' and 'transport as a service' to online consumer service companies. Sam Pan Engineering has it in mind that, if its private cloud is a success internally, then it will offer spare capacity to other companies, as public IaaS. Other opportunities will be to provide value-added cloud services, as ViWi is doing.

Within particular cloud-based business ecosystems there will be many opportunities for specialist service providers. These may include opportunities to provide services that are people-based as well as IT-based, as BPaaS.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

6.4 Quality: improved margin from better service

Increased scale and scope of operation delivers more revenue. Efficiencies resulting from better use of assets and faster operation deliver lower costs. The result is a bigger margin, and more profit.

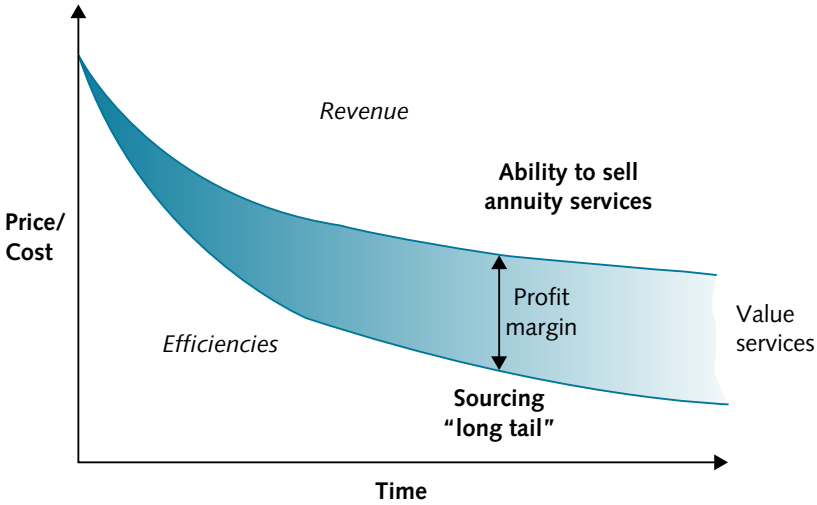

Figure 6.6: Revenue, cost, and margin

6.4.1 Competitive pressure

You will not be alone in adopting cloud computing. Your competitors will adopt it too, bringing pressure to cut price and reduce your precious margin. How do you deal with this challenge?

The positioning of cloud computing, while initially seen as a disruptive technology influence on both buyers and seller prospects, is now evolving into a trade-off between low-cost arbitrage and added value QoS.

In a market driven by cost, there can be intense competition between participants to make cost savings. This is often seen in a negative context, as the lower costs and margins are seen as a detriment to the participants overall. Massively scalable services from cloud computing providers have the effect of driving down costs and prices, as the dynamics of competition are shifted by the presence of potentially rapid cost reductions.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

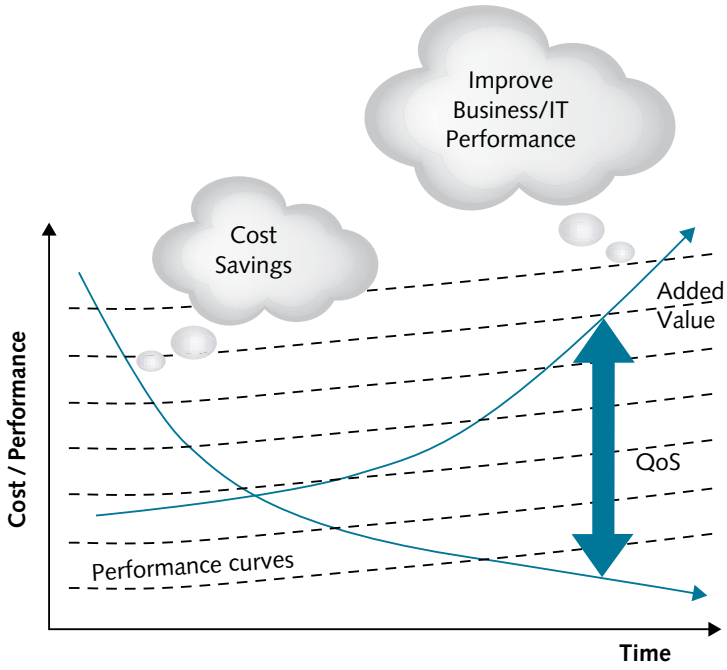

Figure 6.7: Competitive pressure

6.4.2 The importance of quality

The counter-balance to this is the quality of service which, taken in comparison to the cost of that service, characterizes the value delivered. Adopting the cloud computing paradigm drives costs down through greater efficiencies of working capital. Some of the gains must go into efforts to improve quality, if margins are to be maintained.

The differentiator of cloud computing is not just the utility computing services, but includes all the higher-level services that enhance and build service value. This is part of a movement from IT-centric to business-centric services across a wider services continuum, with utility services for infrastructure at one end, and with business-centric software and business processes delivered as a service from the cloud at the other.

6.5 Comparing ROI of cloud and traditional IT solutions

There are a number of things to consider when deciding whether to use cloud computing. Many of them were described in Establishing your Cloud Vision (Chapter 3). They will usually include ROI. If you propose a cloud-based

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

project, you will be asked how the ROI compares with that of a traditional solution. How do you answer this question?

To make the trade-offs between the factors involved, you need a model of how your project will achieve ROI. This model will typically include assets, revenue, workload, costs, time to deliver, features, quality, and scale. It can be in your head, or can be captured by a more or less sophisticated computer modeling tool. Much of the model will be based on normal financial management factors, such as the relation between investment, revenue, costs, and return. Some of these factors will be affected by considerations that are special to cloud computing.

This section illustrates the process, by showing how our three example companies assess the ROI of their cloud solutions. The illustrations use simplified spreadsheet models, which are available for download from The Open Group website [EXAMPLES].

6.5.1 Basis of ROI calculations

ROI is generally defined as the ratio of money gained or lost on an investment to the amount of money invested. (This definition was abbreviated from the Wikipedia entry for ROI.)

When an investor puts money into a venture in expectation of a return, the calculation of the ROI is straightforward. Our third example company, ViWi, provides an illustration of this.

When the term ROI is used in a corporate environment, as with our first two example companies, the situation is more complicated. In principle, ROI can still be calculated as the gain from the investment divided by the cost of the investment, where the gain is the financial improvement resulting from the investment less the cost of the investment, but it is not always easy to determine what the improvement is or what the cost is, particularly when the investment results in an annual profit or saving for the foreseeable future rather than a one-off benefit.

For Konsort Prinz and Sam Pan, we will take the financial improvement to be how much profit or saving is realized over the first three years of the project, and take the cost to be the amount of money required to make the change that enables the profit or saving.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

6.5.2 Konsort-Prinz

The current systems at Konsort-Prinz are not coping with the existing workload, and are limiting the company's ability to do business. The company executives estimate that its turnover and profits would grow by 10% per year if its IT systems were able to handle the increased demand. They have a straightforward decision to make between investing in a major data center upgrade and moving to cloud services. To help them make this decision, they perform comparative ROI calculations for the two proposed solutions. For both calculations, they assume that the costs are incurred during the first year of the project, during which time they continue with their existing IT systems, and that the benefits are realized in years 1 and 2 of the operation of the new system.

They have a budgetary figure of $5M for the data center upgrade, based on quotes received. They take their IT costs of $1.35M for the last financial year as the assumed annual cost of this solution going forward. The cost of operating and maintaining the software application is considered unchanged in both the date center upgrade and the proposed cloud solution. Their calculation is shown in the figure below (Figure 6.8) and gives an ROI over three years of 26%. (These examples ignore cost of capital, inflation, and exchange rate movements. They are important factors in real life, but only serve to confuse the explanations.)

	Total	**Year 0**	**Year 1**	**Year 2**
Increase in profit	6.30	0.00	3.00	3.30
Financial improvement	6.30			
Cost of Upgrade	5.00	5.00		
Investment	5.00			
ROI over 3 years	26%			

Figure 6.8: Konsort-Prinz data center upgrade ROI

The move to cloud services will not entirely do away with up-front payments. The company will need to de-commission its data center, including making severance payments to those IT staff that cannot be re-deployed. They budget $2.3M for this. They will also have some software migration costs. They budget $1.2M for these (again, based on supplier quotes).

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Their cloud service costs are taken from the workload and cost models that they developed as described in Buying Cloud Services (Chapter 4). They have not yet chosen their supplier, but use an annual figure of $550K which seems representative of the three that they have considered. This is based on their current workload, and they increase it by 10% per annum to allow for the increased traffic that they anticipate. The resulting calculation is shown below (in Figure 6.9). The IT costs are shown as a single broad category including specific factors such as audit and security. Similarly, the software migration cost covers all costs associated with the migration, including legal and contract costs. (Note that the numbers are all rounded to two decimal places.)

	Total	**Year 0**	**Year 1**	**Year 2**
Current IT Cost		1.35	1.35	1.35
IT Cost with Cloud Solution		1.35	0.55	0.61
Decrease in IT cost	1.55	0.00	0.80	0.75
Increase in profit	6.30	0.00	3.00	3.30
Financial Improvement	7.85			
Data Center De-commissioning	2.30	2.30		
Software Migration	1.20	1.20		
Investment	3.50			
ROI over 3 years	124%			

Figure 6.9: Konsort-Prinz cloud solution ROI

The calculation shows that the cloud solution will give an ROI over three years of 124% – a very significant improvement over the data center upgrade.

6.5.3 Sam Pan Engineering

Sam Pan Engineering has a comparison between continuing the *status quo* and moving to cloud services. Its corporate executive is considering a proposal to invest $5M in converting about 10% of its servers to private cloud. These servers are massively under-utilized at present, and the proposal assumes a conservative server consolidation ratio of 6:1, so that the cloud servers will do the job of 60% of the servers that the group currently has.

The annual IT costs are currently $25M. The proposal is to keep 40% of the servers, reducing non-cloud IT costs to $10M, and to convert 10% of them

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

to cloud servers, which will then cost $2.5M per year to run. The remaining 50% will be discarded. These will mostly be old machines, with little retained capital value. The proposal includes an investment of $5M to buy and deploy the cloud platform, and for the necessary re-organization of equipment. It assumes that the running costs of the private cloud infrastructure will be the same as the running costs of the current infrastructure. Running a cloud infrastructure is more complicated, but the current costs of managing a silo'ed and not well standardized environment are relatively high, and these factors should balance out.

	Total	Year 0	Year 1	Year 2
Current IT Cost		25.0	25.0	25.0
Non-Cloud IT Cost with Cloud Solution		25.0	10.0	10.0
Cloud IT Cost with Cloud Solution			2.5	2.5
Decrease in IT cost	25.0	0.0	12.5	12.5
Financial Improvement	25.0			
Cost of implementing cloud	5.0	5		
Investment	5.0	5		
ROI over 3 years	400%			

Figure 6.10: Sam Pan Engineering ROI calculation

The calculations show payback of the investment less than five months into the second year, with ROI over three years of 400%.

6.5.4 ViWi

ViWi is the only one of our three example companies to which the term 'return on investment' applies in the literal sense. The company has backers who are investing capital, and want to know what return they will get.

ViWi's founder is asking for $400K which he believes will enable him to establish the company. The backers are looking to invest money over a period of one or two years.

ViWi's founder does some calculations, produces a nice graph showing revenue and costs, and says that the company will be in profit within a year, and over two years will have an accumulated surplus of revenue less cost of nearly $5M.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

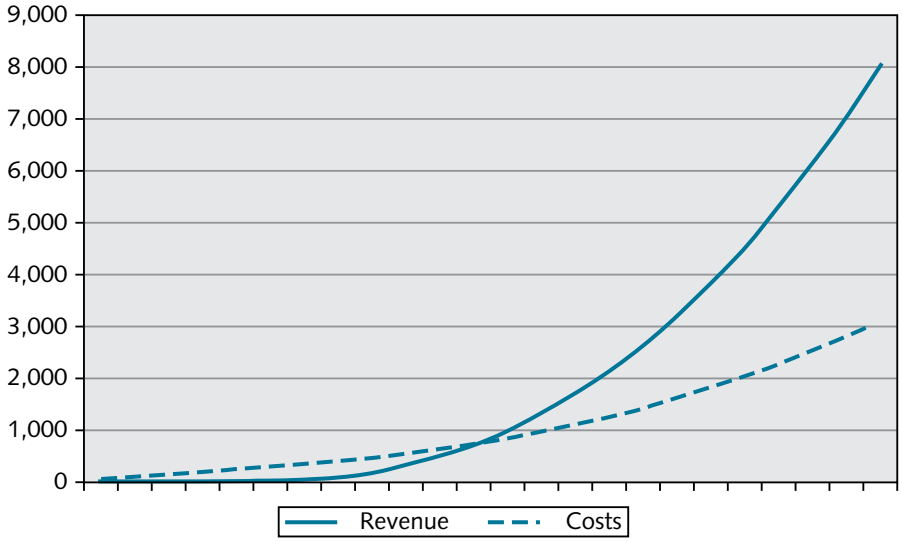

Figure 6.11: ViWi Years 1 and 2 cumulative revenue and costs

Under cross-examination, he explains that this is based on an underlying assumption of 50% per month compound growth in virtual widget use. They ask what if growth is only 10% per month? He admits that this would be a loss-making situation, but points out that 100% per month is perfectly possible also, and this would give a surplus of nearly $12M, just in the first year!

The fact is that this is a very risky investment. The backers must use their business judgement as to the likely take-up of the product, and to assess the risk involved.

They agree to pay $400K for a 50% share in the company. If the value of the company after two years is $3M (which would be roughly the amount of cash generated, after tax, under the founder's projection), and their share is therefore $1.5M, the ROI over two years would be 375%, as shown below (in Figure 6.12).

Profit after tax	3.00
Amount gained by investors	1.50
Investment	0.40
ROI Over 2 Years	375%

Figure 6.12: ROI for the backers of ViWi

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

The interesting point from the cloud computing perspective is that the risk is in the normal business factors, not in the use of cloud computing. The ROI figures are affected very little by changes in the assumptions made in the cloud workload and cost models, because the costs are directly related to the workload, and hence to the revenue. And there is no capital cost for buying development and deployment systems, which would add to the loss considerably if the product was a failure.

6.6 Measuring and tracking ROI

Clearly the normal financial models showing revenue, capital and operational expenditure, and costs are the main models from which ROI is calculated. For cloud projects, you should also look at indicators of health of the four ROI drivers described earlier in this chapter: utilization, time compression, scale, and quality. These are often the leading indicators of project ROI.

6.6.1 Utilization

Utilization is a key factor that affects productivity, as explained under Productivity: More Business with Less IT (Section 6.1).

At the infrastructure level, the *instance to asset ratio* shows the number of guest operating system instances to physical resource assets. This is a key metric for IaaS providers.

The *tenancy to instance* ratio measures the number of tenants per resource – see Resource Pooling (in Section 1.2). This ratio is important to providers of PaaS and SaaS. For PaaS, it is the number of separate platform environments per platform instance. For SaaS, it is the number of client partitions per application instance.

CPU utilization and *memory utilization* are important to providers and to IaaS and PaaS consumers. Consumers, PaaS providers, and SaaS providers should look at them in conjunction with *workload*. If CPU and memory utilization are low then cost reductions can be achieved by allocating more workload to each resource.

Consumers should of course monitor *cost*, and should look at it in conjunction with workload. If payment is directly related to resource use –

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

'pay-by-the-drink' – then costs should track workload. If there is an element of time-based payment – 'pay-as-you-go' – then this may not be the case, and there may be scope to lower costs by tuning resource usage.

Private cloud users often look at the *server consolidation ratio* – the ratio of servers in their private cloud to those that were used before the cloud was introduced.

6.6.2 Time compression

These metrics relate to the factors discussed under Speed: Getting There More Quickly (Section 6.2).

For consumers of cloud services, *speed of provisioning* is an important metric. As mentioned earlier, speed of provisioning is crucial for ViWi.

Speed of provisioning contributes to *speed of cost reduction* (see Figure 6.5), which can provide a long-term indicator of success.

6.6.3 Scale

Scale of operation is discussed under Size: Breaking New Ground (Section 6.3).

Workload is probably the best IT measure of this. It should be looked at in conjunction with *revenue* and *cost*.

Cloud providers – including providers of private cloud – will also monitor their *investment in IT resources*.

You may be able to measure *market disruption rate* as a combination of rate of revenue growth and rate of new market acquisition.

A measure of the complexity of your ecosystem interactions is *ecosystem optionality* – the number of external assets, APIs, and services that you have to deal with.

6.6.4 Quality

The importance of quality is explained under Quality: Improved Margin from Better Service (Section 6.4).

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Basic QoS can be measured in terms of *availability*, *reliability*, *recoverability*, *responsiveness*, and *throughput*. You may also wish to monitor *manageability* and *security*. These are all described for services that you consume under Establishing Requirements (in Section 4.2); the descriptions apply equally to services that you provide, which is what we are concerned with here.

You should monitor *user satisfaction*. You cannot rely on your own assessment of the quality that you deliver. It is best to do this proactively. There is a well-known and successful Voice-over-IP supplier that, at the end of each call, gives the user the opportunity to comment on call quality through a simple web form. By contrast there is another provider that, though excellent in many respects, requires complaints on quality to be made to its help desk – but there are many competing providers, it is easier to switch to a different one than to complain, and that is what some users are doing.

Customer retention is a useful indicator of quality. ViWi plans to keep close track of whether its customers continue to use the virtual widgets over a long period, or whether they forget about them after the first few days of interest.

Revenue efficiencies, measured as a combination of margin increase per unit revenue and rate of increase of 'annuity' income, can also help you to determine whether you are succeeding in your attempt to respond to competitive pressure by raising quality.

6.6.5 Examples

To illustrate the principles of monitoring cloud ROI, let's look at how each of our example companies takes stock of its ROI position part-way into its project. To keep things simple, the examples just show indicators for utilization, speed, scale, and quality, mostly with a single indicator for each. The indicators are measured using a scorecard approach on a scale of unacceptable/poor/average/good/excellent.

6.6.6 Konsort-Prinz

The indicators are:

- Utilization – the ratio of used processing power to provisioned processing power
- Speed – the time taken to provision new customers and suppliers in the system

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

- Scale – the total volume of workload that can be handled in a year
- Quality – user satisfaction, measured through a program of questionnaires and interviews with staff

These are compared with the values of the in-house system, which is the benchmark for 'average'. Konsort-Prinz expects an improvement in utilization, and also in scale due to the ability to provision cloud resource, but no significant change in the other figures.

In the third week of November of the first year of operation of the new system, there is a crisis executive meeting. The operations manager says that the system is no good, they are not coping, they will not be able to handle the Christmas rush of orders, there is nothing they can do about it now, but they must scrap the cloud project and invest in a new and bigger in-house system for next year.

The financial implications of this are frightening. The company will make a loss for the year, and lose customers. It will have to borrow heavily to develop a new system – which will be difficult in view of its worsening financial position. It might even go bust.

The ROI scorecard shows that its quality level is unsatisfactory, but the other measures are healthy.

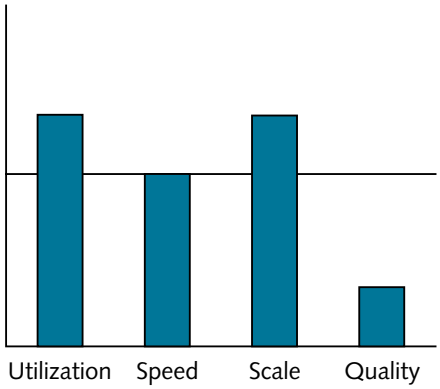

Figure 6.13: Konsort-Prinz November ROI scorecard

Utilization is good – and costs are actually marginally lower than forecast. Time taken to provision new customers and suppliers is as expected. The

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

system has not reached any CPU or memory limits, and has plenty of capacity to handle more business.

The main complaint is about response time. The system is very sluggish, and it is taking longer than it should to process orders. The staff who would normally be able to handle the December peak can hardly manage the November workload.

Fortunately, there is a solution. The utilization target was too aggressive. The crisis meeting decides to allow a higher level of resource provisioning to improve response time. This increases costs, but not dangerously so – and the December peak is then handled without any problems.

By taking advantage of cloud scalability, the company has been able to fine-tune its operation and avoid crisis measures.

6.6.7 Sam Pan Engineering

The indicators are:

- Utilization – server consolidation ratio
- Speed – the time taken to convert a server to become part of the private cloud
- Scale – the expected proportion of projects using the private cloud
- Quality – user company satisfaction

Six months into the project, there is no crisis, but the ROI scorecard does give the project team some cause for concern.

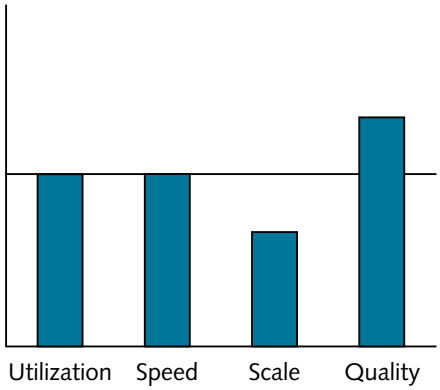

Figure 6.14: Sam Pan Engineering six-month ROI scorecard

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Utilization is about as expected, in fact slightly better, with a server consolidation ratio of 8:1. Server conversion time is as planned. The quality indicator is good; despite initial reluctance, the companies that are using the new system are finding that it works well for them. There is even one high-profile water-purification project where ready availability of cloud resource enabled the company to complete some crucial modeling work ahead of schedule. But the scale indicator is disappointing. The constituent companies of the group are resisting the adoption of the new cloud system, and it now looks as though only half of all projects will use it.

Assuming that the investment required will be only half of what was planned, there is no change to the percentage ROI over three years, but the financial improvement is only half of what corporate management was expecting. This is shown below (in Figure 6.15). (Note that the numbers are rounded to two decimal places.)

	Total	**Year 0**	**Year 1**	**Year 2**
Current IT Cost		25.00	25.00	25.00
Non-Cloud IT Cost with Cloud Solution		25.00	17.50	17.50
Cloud IT Cost with Cloud Solution			12.5	12.5
Decrease in IT cost	12.50	0.00	6.25	6.25
Financial Improvement	12.50			
Cost of implementing cloud	5.00	5.00		
Investment	5.00	5.00		
ROI over 3 years	150%			

Figure 6.15: Sam Pan Engineering reduced scale ROI projection

The project team discusses the situation with corporate management. Because the other indicators are good, the project leaders are able to make the case for exerting more pressure on constituent companies to adopt the new system. Corporate management decides to publicize the success stories more strongly, and to implement budget cuts for companies wishing to keep their in-house IT. This combination of measures brings about the desired improvement in take-up.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

6.6.8 ViWi

The indicators are:

- Utilization – the resource cost per widget
- Speed – the compound monthly growth rate in the virtual widget population
- Scale – the expected eventual widget population size
- Quality – the proportion of active widgets and the subscription renewal rate

Six months into the project, there is good news and bad news.

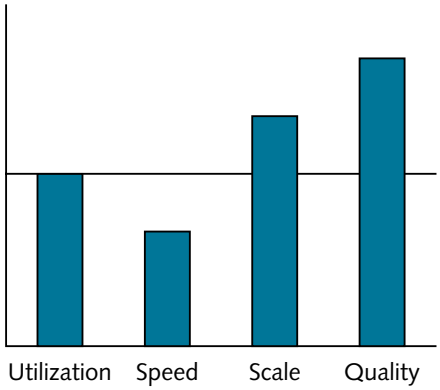

Figure 6.16: ViWi six-month ROI scorecard

The good news is that the company has been able to launch after three months, rather than five months as anticipated. The assumptions about time-to-market, partly based on experience with traditional developments, has proved pessimistic. Also, the scale and quality indicators are looking good. The bad news is that the growth rate in the widget population is only 25%, half of the 50% rate that was assumed. The impact of this is that the project will make a loss of about $360K in the first year, but hopefully show a profit of $200K after two years, giving a ROI of 15%.

Profit after tax	0.12
Amount gained by investors	0.06
Investment	0.40
ROI Over 2 Years	15%

Figure 6.17: ViWi ROI projection after six months

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

The backers have put in $275K so far. An additional $85K will be needed to carry it through its loss-making period. Should they can the project?

Were it not for the early launch, the answer might well be "Yes!" But the rapid progress has given confidence and, more importantly, has given some time to show the value of the product. The average proportion of active widgets is 11% as compared with the 5% projected, there have been few subscription cancellations so far, and the widgets are being used in ways that had not been expected, indicating a larger potential market.

The backers agree to continue the project. More than that, they will invest extra money in advertising, to try to increase that growth rate.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Chapter 7

The challenge

We have looked at why people are turning to cloud computing, at what it is, at how it is used, and how to establish a cloud computing vision. We have discussed the issues involved in buying cloud services. We have gained an understanding of the risks of cloud computing, and of how to build and measure its rewards.

Used in the right way, cloud computing can be a great accelerator of collaboration-oriented architecture. But what does that mean for the enterprise itself, and for its executives and decision-makers?

A common point of discussion pertaining to cloud is the scope and role of the enterprise. In fact some have proposed to consider changing the viewpoint of 'cloud in the enterprise' to 'enterprises in the cloud'. This is a subtle yet significant change that presents a transition in thinking and enterprise operations, from silo or boundary management to 'boundaryless' services.

Cloud computing is a change brought about by the convergence of a number of new and existing technologies. It is an important stage in the development of IT systems, comparable with the emergence of the mainframe, the minicomputer, the microprocessor, and the Internet.

This is happening in the context of major transformations in business and society as a whole, with which cloud computing is inextricably linked. The trends are described in books such as Why Buy the Cow? from Subrah S. Iyar on how on-demand powers the economy [IYAR] and The Big Switch by Nicolas Carr that takes the analogy even further with distributed industrial IT services [SWITCH]. We are apparently headed for virtual businesses, serving a virtual society.

Economic, environmental, and global activities shape regional markets, products, and services in many industry sectors. Government legislation

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

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-based networking has become part of everyday social interaction. Its impact on business could be equally profound. Collaboration services are now just beginning to be used in commercial companies. They will become an essential part of our business fabric as the children of today, schooled in their use, become the business leaders of tomorrow.

The business fabric is changing in other ways. Business units are becoming more specialized, and therefore more efficient. Each unit operates within a set of ecosystems, performing particular services for the other participants, and using services that they provide. The enterprise is one of those ecosystems. The bonds that bind its constituent units together are becoming weaker, and their external bonds are becoming stronger.

Cloud computing is an enabling factor. The Internet enabled business units to communicate faster and more effectively. Cloud computing enables them to share software applications that embody business logic, transforming the speed and efficiency of ecosystem operation.

Cloud is an Internet phenomenon. The Internet way of doing business is very different from traditional enterprise models. Enterprises that try to base their cloud initiatives on models where everything (processes as well as IT) is internal, will never achieve real returns from the cloud. In short, the enterprise architecture has to change to embrace the new way of doing business. The change will not be overnight but a start needs to be made now.

The move to cloud has a particular impact on the IT department. Traditionally, it was wholly within the enterprise and served other enterprise business units. It met all of their IT needs, designing, procuring, and operating the IT infrastructure and the business applications that ran on it. With cloud computing, some or all of those services are performed outside the enterprise. Cloud computing is extending the trend to remove services from IT-department control that started with the introduction of the PC. The business of IT becomes that of leveraging the products and services through

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

competing platforms and channels to defend, attack, and build customers and market share.

Most enterprises apply cloud using an incremental approach. They may start by incorporating external cloud services in their in-house solutions, and gradually move to buying whole solutions such as CRM as external services. Incremental adaptation is less risky than big-bang. Taking things one step at a time helps to ensure that the components of your architecture remain interoperable.

Cloud does not just affect the IT department. It impacts on the whole business:

- It changes what the business does. This is not just a question of whether the enterprise does its own IT. It transforms the way that business units interwork with other participants in their ecosystems, and this may lead to the whole shape of the enterprise being redefined.
- By transforming the way business units within the enterprise interwork, it can change the internal structure of the business too.
- It changes what the people in the enterprise do. Internal business processes have already become significantly computerized. Collaborative business processes will become computerized in the same way.
- It changes the way that people work. Social networking tools will change business, as well as social, interaction. One change will be to where people work – they will have less need to share an office, and more opportunity to work from home; cloud computing is an enabler for 'the new way of working', and the desire to work this way is a driver for cloud adoption. Other changes will be to improve the quality of what people do, and increase their productivity.

There are many implications of using cloud for business and IT stakeholders. The results can be very positive. But this is not guaranteed. Spending on cloud services could go out of control. Employees could refuse to cooperate in the transition process. The cloud services may not integrate with the business model, or transform it in the way that was intended or desired.

To take advantage of cloud computing, you must develop a vision for the future of your enterprise, and transform the enterprise to realize that vision.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

The vision may start with new IT services and new ways in which they will be provided. It may go on to define a new shape for the enterprise, and a new internal structure. It may include new principles of operation and new processes. It may mean a changed way of working, and a new culture.

This will mean a major architectural transformation, affecting the business processes and the supporting information systems and technology. It should be a positive experience for the enterprise, the people that work in it, and yourself. It will require careful consideration of the factors, discussion with all the people concerned, and big, sometimes painful, decisions.

This is the challenge that cloud computing gives to the business executive.

Are you ready?

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Appendix A

Cloud computing in use

This appendix contains a set of use-cases or scenarios, which provide a high-level depiction of real-world business problems that were addressed using cloud computing. It also includes analyses of the actors involved in the use-cases and of the benefits of cloud computing that the use-cases exemplify.

The use-cases were developed by the Cloud Business Use-Cases project of The Open Group Cloud Computing Work Group. The project worked in co-operation with other cloud computing use-case industry initiatives, but focused specifically on business use of the cloud as opposed to its technical operation. Its findings are published as an Open Group White Paper [CBUC].

The project developed 24 use-cases. Each case considered a typical use of cloud computing by a business. The uses concerned were in real businesses, known to or researched by the project team members, but sometimes represent abstractions of the experiences of one or more enterprises rather than a record of the experience of a single enterprise.

The use of cloud computing is increasingly widespread. The project team found that the overall rate of take-up across industries is pretty consistent, although the use of private and/or community clouds is more common in some specific industries.

The use-case descriptions cover the following aspects:

- Category – industry, sector, or environment
- Company background (not in all cases)
- Business problem/description of the situation
- Actors impacted by the business problem
- Key business requirements (including quality and/or service-level requirements where known)
- Business risks associated with the problem

Six of these use-cases are cross-industry and 18 are industry-aligned.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

A1 Cross-industry use-cases

Case 1: Business activity reporting

Company Background	This cross-industry company must address the need to produce different reports for Key Performance Indicators (KPI) and Service-Level Agreements (SLA) to plan, understand, and optimize operational and financial performance.
Business Problem/ Description	The reports/status information for numerous stakeholders is required in varying formats. However, data for creating reports is in disparate locations, and the data needs to be aggregated, filtered, and assembled in order to be more useful and generate value for the target audience.
Actors	Shareholders Operational Management Service Providers
Business Requirements	Need to generate management reports for shareholders using the same underlying data, but in a different format, as those used for operational management to report the business performance to shareholders – this can be a powerful way of driving efficiency and complying with exchange listing requirements. Need SLA/OLA reports for service providers so that various internal and external service providers can have their performance reports produced from disparate locations and organizational entities.
Business Risks	Inadequate governance due to lack of timely and inefficient reporting Lost customers if SLAs or OLAs not met Loss of transparency for stakeholders

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Case 2: Modernization

Company Background	This cross-industry company must address modernization of its business processes and legacy applications. However, it does not have the skills in-house to conduct such activities. Existing business services and assets are out-of-date and need modernization.
Business Problem/ Description	Business needs to combat the problem of 'legacy-ization' – it needs to change existing assets for new assets and skills rapidly. Business needs access to new technology quickly – the innovator's dilemma [CHRISTENSEN]. Business does not have the skills in-house to transition via cloud facilitation to modernize business processes.
Actors	Business Operations Controller Business Procurement Management Innovation Management Portfolio Management Data Center Management IT Operations Management
Business Requirements	Need to modernize infrastructure Need to modernize applications Need to modernize business processes Improve employee skills
Business Risks	Lost competitiveness Lost opportunities due to inability to support new business market Lost skills and resources retention Excessive investment spends on maintaining legacy applications Impact on existing contracts for support, licensing, and services

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Case 3: Rapid business capacity and scale

Company Background	This business needs to scale up its operations rapidly, including increased IT capacity – within a short cycle of days to a few months – to meet specific operating workloads.
Business Problem/ Description	There is deficiency of resources and capacity to meet business activity demand to support day and night-time peak loads. The company is facing issues such as how to balance compute workloads better, how to optimize costs of operations, and how to follow variable demands of service effectively and efficiently at lower cost.
Actors	Business Operations Controller Business Procurement Management IT Operations Management
Business Requirements	Remote user desktop services Infinite compute and storage capacity on-demand Access to temporary computing processing for business simulation and transaction processing Business process extensions; e.g., BPO Access to rapid deployment of existing business process resources including applications and infrastructure
Business Risks	Lost revenue from lack of capacity Lost customers from under-performance of business delivery

Case 4: Operational efficiency

Company Background	A company embarked on a strategic initiative to improve its management of development and test infrastructures.
Business Problem/ Description	Business requires better management of development and test infrastructure without increasing IT labor and set-up costs.
Actors	Application Development and Delivery Test and Quality Assurance Business Operations Controller Business Procurement Management IT Operations Management

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Business Requirements	Provide management (e.g., automatically or self-service) for business activities that scale up and down operations with associated temporary capacity Provide governance structure and control to manage disparate activities for project-based resources
Business Risks	Opportunities for analysis missed Products and services are late to market Lost customers and market share due to failure to create ease of doing business Excessive complexity in business operations

Case 5: People productivity, workplace collaboration

Company Background	A company embarked on a strategic initiative to improve its internal and external communications.
Business Problem/ Description	The company's workforce needs to communicate and collaborate more efficiently within the company and with its global network of customers, suppliers, and business partners.
Actors	Operations Management CIO CTO
Business Requirements	Remote user desktop services Improve the internal collaboration requirements (e.g., communication methods) for the business Improve the external collaboration requirements (e.g., customer meetings) for the business Improvements must be easy to implement Improvements must provide secure connections
Business Risks	Improvement implementations costly and time consuming Improvement connections not secure

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Case 6: Development and test services

Company Background	A company embarked on a strategic initiative to establish development and test environments for functionality and performance of applications so that businesses can achieve faster time-to-market for their products.
Business Problem/ Description	The development and test environments need to be provisioned at short notice and then 'mothballed' or redeployed elsewhere once the development or tests have been conducted. This represents a significant cost in effort and capital expenditure (CAPEX). Virtualization can be used to address this but there is often a demarcation along physical lines across business units, and so the take-up rate is not as large as it should be due to the localized nature of using virtualization.
Actors	Development Teams Test Teams Product Sales and Marketing Executives
Business Requirements	Development and test environments need to be provisioned at short notice. Development and test environments need to be 'mothballed' or redeployed elsewhere once tests have been conducted. Time-to-market for products needs to be maintained or decreased.
Business Risks	Time-to-market increases CAPEX costs increase as individual projects or business units need to provision test environments that are segregated

A2 Financial services use-cases

Case 7: Business transaction assurance and continuity

Company Background	This financial services company is seeking rapid continuity services that assure business operations and associated transactions, particularly during mergers and acquisitions.
Business Problem/ Description	This company lacks the necessary in-house skills and resources to support the high degree of sensitive assets and data that is required to support its key business processes. Business Transaction Assurance, Disaster Recovery Planning, and Testing of vital business processes for the same are required to comply with enterprise and regulatory guidelines.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Actors	Business Operations Controller Business Procurement Management IT Operations Management
Business Requirements	Provide necessary disaster and recovery back-up processes Ensure that business activities are met during times of additional complexity, such as mergers and acquisitions Access to rapid deployment of compatible and simple end-user services through browser or end-user device
Business Risks	Current business operation is exposed to disaster and recovery back-up processes either being insufficient or not in place Lost information contributing to poor business decisions Loss of business continuity across processes and new ventures Significant loss of assets or revenue flow

Case 8: Business market launch

Company Background	This financial services company needs to rapidly introduce new products and/or services into a new or existing marketplace.
Business Problem/ Description	There are insufficient resources and capacity for this company to rapidly respond to marketplace needs including seasonal events, although new entrant opportunities have been identified.
Actors	Marketing Management New Product Development Business Operations Controller Business Procurement Management IT Operations Management
Business Requirements	New products and/or services require new applications Rapid development and launch of new product and/or service specification Reliable and responsive service provision to support product and service launch Need more capacity for storage of information and to support service delivery
Business Risks	Lost opportunity cost Lost customers from under-performance of business delivery

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Case 9: Secure business assurance services

Company Background	This financial services company is seeking transaction verification and hosting security assurance services.
Business Problem/ Description	Security assurance is paramount to business operations, yet there is increased concern over the external service provider's services. This business lacks internal skills and resources to develop security services. This business seeks to gain cloud computing elasticity benefits but it is unable to support public service access due to security restrictions. Business legislation requires the company to use its own private secure services to control access to secure business services.
Actors	Business Operations Controller Business Procurement Management Security and Compliance Data Center Management IT Operations Management
Business Requirements	Private cloud capability Access, authentication, repudiation Audit and e-Discovery compliance Securely partitioned and isolated Disaster recovery Web security authentication Business continuity
Business Risks	Security and compliance breach Lost business continuity Sovereignty compromise

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

A3 Government use-cases

Case 10: Research

Company Background	A US Department of Energy research and engineering facility decided to replace a Cray XMP accessed via the Internet with a cluster of UNIX servers. This decision was made both on the basis of cost (upgrade *versus* replace), and a goal of advancing clustering technology.
Business Problem/ Description	Researchers and their collaborators were demanding more compute capacity than was available on the existing Cray XMP. The business problem was to allocate compute resources equitably across a mixture of batch and interactive workloads, and to bill each project for the resources used. 'Controlled anarchy' was the researchers' term for the acceptable level of systems management.
Actors	Business Operations Controller Business Procurement Management Mobile End User IT Operations Management
Business Requirements	Equitably allocate the compute resources across a mixture of batch and interactive workloads Bill each project for its use of compute resources Support the then common programming languages to achieve minimum processing performance for typical high-performance computing workloads Allow users to submit distributed workloads Optimize resource utilization Keep each project's usage within its budget Scale capacity quickly and cheaply
Business Risks	Systems management tools and programming models were not sufficiently mature to keep the IT costs within budget.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Case 11: Productivity

Company Background	This government agency wants to improve employee productivity and reduce end-user costs by leveraging mobile and social communities to improve the quality of business information and knowledge exchange.
Business Problem/ Description	This entity must take advantage of lower-cost IT options including rapid delivery of productivity tools to its workforce. The organization contains many mobile workers who regularly need access to services; however, it is unable to invest in dedicated (in-house) mobile services and support. It is also experiencing huge desktop costs and support due to a large user base with many locations and disparate business operations.
Actors	Business Operations Controller Business Procurement Management Mobile End User IT Operations Management
Business Requirements	Business mobility services and support Business productivity requirements to improve community and knowledge development in the organization Reduce end-user costs/charge-back to customers
Business Risks	Inefficiencies in business productivity Technology incompatibility contributing to poor systems performance and downtimes Security issues introduced with remote workforce Dissatisfied customers

Case 12: Business utility/agility

Company Background	Current business unit activity has opportunities to create common utility services to better support merger and acquisition activities.
Business Problem/ Description	This government agency is prone to expanded operations, mergers, and/or acquisitions, and activities that are shrinking the capacity of its IT.
Actors	Program Managers Business Operations Controller Business Procurement Management IT Operations Management

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Business Requirements	Use commodity services to reduce operating costs Rationalize business assets through alternative sourcing of assets Use external transition services to assist consolidation and rationalization of merger and acquisition activity Need more elasticity – access to and rapid deployment of end-user services
Business Risks	Over and under-provisioned capacity Failure to leverage assets to better serve common business services across business agency units Complex asset dependencies remain unresolved Lost customers from underperforming business Excessive costs of infrastructure assets

Case 13: IT cost optimization

Company Background	Agency is conducting activities that are affecting IT capacity. These activities may include operations' expansion, divestment, or significant mergers and/or acquisitions. The business needs to manage the cost of business operations and seek alternative ways to optimize investments.
Business Problem/ Description	Current costs of IT are too expensive – in addition, productivity tools for the workforce are not delivered fast enough to keep up with market demands.
Actors	CTO, CIO, COO Business Operations Controller Business Procurement Management Order Fulfilment, including Storage System Administrators IT Delivery
Business Requirements	Current investments in business are high and need to be rationalized The business seeks supply cost reductions through alternative methods of service supply on common services Current business unit activity has opportunities to create common utility services to support merger and acquisition activities Accurate, consumption-based charges for services
Business Risks	Complex asset dependencies remain unresolved Excessive costs of infrastructure assets including storage Failure to leverage assets efficiently in order to provide common business services across business partners and enterprise domains

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Case 14: End-user cost optimization

Company Background	Business needs to facilitate delivery of new business services to users to improve productivity. Business wants to reduce costs of services with rapid deployment of mobile devices and create better social community experiences to improve access and quality of business information along with improved knowledge exchange.
Business Problem/ Description	Current business operation is looking to maximize low-cost IT options to reduce costs and rapidly deliver business productivity tools to its workforce. It has many mobile workers who regularly need access to services on the move, but lacks the scale to invest in dedicated mobile services and support. In addition, there are huge costs in desktop and support due to a large user base, diverse locations, and disparate business operations.
Actors	CTO, CIO, COO Mobile Users Delivery Teams, including web and application hosting Business Partners
Business Requirements	Provide services and support on-demand to mobile workers Improve the community and knowledge development in the organization Need for large-scale investments in end-user services with high costs and complexity
Business Risks	Business productivity and efficiency decreases Business loses customers from under-performance of business delivery Security issues due to workforce working remotely Incompatible business processes due to IT insertion IT complexity increases due to additions to new business process capabilities

Case 15: Business growth, development, test, and validation support

Company Background	Business unit has very elastic processes that require rapid test, validation, and quick solution deployments.
Business Problem/ Description	Business has rapid introduction of new products and services with seasonal variation and on-demand needs for assets and resources, yet it faces budget constraints and shortages in IT capacity.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Actors	Business Operations Controller Business Procurement Management IT Operations Management Test Management Data Center Management Development Management
Business Requirements	On-demand test service On-demand business change validation Access to development and test services
Business Risks	Over budget on development and test activities Failure to deliver business products and services Lost customers and market share due to failure to deliver customer promises and needs

A4 Telecommunications operator use-cases

Case 16: Service delivery

Company Background	Current business is seeking online shopping and consumer services to sell products and services through mobile, cell, and other network-based services.
Business Problem/ Description	This business is seeking to move services to new markets and channels through network-based services, but managed network service enablement is currently not available. The communications network is seen as a challenge for extending products and services to a wider range of customers and markets that access network-based channels.
Actors	Business Operations Controller Business Procurement Management Internet Service Provider (ISP)/Network Service Provider (NSP) Data Center Management IT Operations Management
Business Requirements	Development of self-service and consumer market services to support product delivery and development Transport as a service Network as a service Access to rapid deployment of compatible and simple end-user services through browser or end-user device
Business Risks	Lost market opportunities due to inefficient communications and delivery Network performance issues Lack of experience in managing network-based services

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Case 17: Partner collaboration

Company Background	Agency needs to facilitate collaboration with a range of partners and customers with no or limited IT resources or funding. There is a large marketplace of individual customers and participants all with common service needs to access business offerings.
Business Problem/ Description	High manual workload tasks make each business transaction increasingly less cost-effective with increasing volumes. Business needs to offer common services to a large market that involve interaction of participants.
Actors	Business Procurement Data Warehouse Management Data Center Operations
Business Requirements	The business needs to manage the cost of business operations and seek alternative ways to offer common services Provide social collaboration tools and skills within the business to meet demand Information and services need sharing between participants' collaboration using a common collaboration process and tool set Control transaction costs
Business Risks	Transaction costs increase Collaboration is not efficient Customers' and participants' business is lost

A5 Media and entertainment use-cases

Case 18: Test assurance services

Company Background	Company has many test and validation procedures to ensure business operations. Test assurance services of business processes need accessing to ensure validation and continuity of services and changes to services. Business introduces many new products and services – continually updating the business portfolio.
Business Problem/ Description	Company has rapid introduction of new products and services with seasonal variation on-demand for assets and resources, yet it is often constrained in capacity.
Actors	Media Outlets Online Shoppers Data Center Operations

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Business Requirements	Business has very elastic processes that need rapid test and validation Business needs to improve costs of complex development and test overhead
Business Risks	Development and test budgets exceeded Business products, services, and delivery are constantly changing Delivery of business products and services fail Lost customers and market share due to failure to deliver customer needs

A6 Health services use-cases

Case 19: Rapid deployment

Company Background	Business needs to facilitate new business services to business users to improve productivity experience. Business wants to reduce costs of services with rapid deployment to mobile devices. Business wants to create social communities to improve access and quality of business information and knowledge exchange.
Business Problem/ Description	Current costs of IT options too expensive and productivity tools for workforce not delivered fast enough.
Actors	Healthcare Providers Healthcare Administrators HealthCare Payer/Payee Patients IT Operations
Business Requirements	Operations looking to maximize low-cost IT options to reduce costs and rapidly deliver business productivity tools to its workforce Current business has many mobile workers who regularly need access to services on-the-move, but lacks the scale to invest in dedicated mobile services and support Huge costs in desktop costs and support due to large user base with many locations and disparate business operations
Business Risks	Lost business productivity and efficiency Lost customers from under-performance of business delivery Security issues of workforce working remotely Risk of compatibility of IT technology insertion Increased complexity of IT with additions to new business process capabilities

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

A7 Pharmaceuticals use-cases

Case 20: Rapid business process extension

Company Background	Business needs to introduce a new business process. The solution can be an extension to the existing applications through use of rapid BPM or a sourced solution from a cloud supplier or service source forge inventory. Need to rapidly access IT capability not available in current business operations. The solution seeks to source a business process function available from the open market.
Business Problem/ Description	The company is currently operating in silos, which are causing disconnection between how business and IT resources are shared and coordinated. This is negatively impacting manufacturing and other key business processes. Development and test resources appear to have a deficit in skills to support quality delivery of the new business process.
Actors	Business Operations Controller Business Procurement Management Market Sourcing and Procurement IT Operations Management End Users/Patients/Healthcare Administrators
Business Requirements	Business needs to introduce a workflow to connect different business processes and applications Business needs to introduce a workflow to connect different business applications to improve information quality flow Need to introduce a temporary project administration Need to introduce a temporary development and test environment to develop and deliver a new functional requirement Need to introduce quality management tracking and reporting of business processes cutting across a number of business areas and applications and database sources
Business Risks	Lost customers due to under-performance of business delivery Risk of compatibility of technology insertions Increased complexity of IT Poor quality products and services Increased complexity of IT with additions to new business process capabilities

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

A8 Distribution use-cases

Case 21: Brand unification

Company Background	A retail distributor embarked on a strategic initiative to unify its brands worldwide under the company's name for competitive advantage, and to increase business results by "helping its multiple business units work together more efficiently".
Business Problem/ Description	The company's workforce needs to communicate and collaborate more efficiently with its global network of customers, suppliers, and business partners in support of a strategic brand unification initiative.
Actors	Suppliers Line of Business (LoB) Leaders Distribution Centers
Business Requirements	Reduce total costs of ownership by: transforming the workforce into a globally integrated team; looking to partners for solutions that may have better capabilities; and leveraging existing IT infrastructure without burdening IT resources Be innovative without compromising security, reliability, or privacy
Business Risks	Unmet business performance objectives Poor partner selection and product service offerings Delayed order fulfilment cycles Inconsistent service offerings Security, reliability, and privacy compromised

Case 22: Innovation – new ventures

Company Background	This distribution company is small to medium-sized with a broad and diverse customer base. This was a new, innovative undertaking.
Business Problem/ Description	The company needed a business model that was responsive and flexible to diverse customer demands. Current processes were reported as co-location with personal management of purchasing, configuring, and ongoing maintenance of hardware, software, and operating systems. The process was time-consuming, costly, and appears labor-intensive. The company needed provisioning and de-provisioning of servers in minutes, instead of its current turnaround rate of days.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Actors	Internet/Web Users COO & Operations Staff Data Center Staff & Management
Business Requirements	The company was challenged to scale its business in a cost-effective and timely manner in order to achieve customer satisfaction targets that optimized costs to the customer and the company as a whole Two business performance indicators were to: provide spam filtering at a low price to consumers with little or no venture capital funding; and achieve 99% spam blocking rates with continued operations as a cost-effective company
Business Risks	Data center requirements were reported as 'in flux' and at risk in manageability and implementation. This caused concern regarding the company's overall support capabilities of this new undertaking.

A9 Energy and utilities use-cases

Case 23: Green footprint

Company Background	Business needs to facilitate all options to improve 'green footprint' emissions of their operations. Currently, business has expectations for rapid growth in IT services demand causing 'drift' in its workloads and emissions from increased power consumption.
Business Problem/ Description	Current business operation investments internally have reached limits in green cost reductions and are looking for other options and sources.
Actors	Data Center Management and Operations Line of Business (LoB) Leaders Sustainability Sales and Marketing Executives
Business Requirements	Maximize utilization Reduce carbon footprint
Business Risks	Lack of carbon reduction compliance Lost image in marketplace over green credentials Lost business due to inability to offer green services

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

A10 Higher education use-cases

Case 24: Virtual learning

Company Background	Distance learning is an essential component of public community colleges that are engines for workforce development for local economy. In recent years, public community colleges are adversely affected by state budget cuts. These college campuses are looking for opportunities to cut cost without affecting delivery channels and student enrollment.
Business Problem/ Description	Virtual learning environment is one of the common delivery methods in higher education. Each of the community colleges in a state maintains its own set of applications and infrastructure for virtual learning environment duplicating IT cost for each of the campuses and college systems.
Actors	Faculty Student Instructional Technology Specialists
Business Requirements	Use of shared environments by different community college campuses in a state for delivery of virtual education to certain population in the community as an alternative to classroom teaching
Business Risks	Lack of sufficient state funding will lead to elimination of distant learning programs Shutting down a delivery method which is often the only viable learning delivery method to a part of the population who are engines for the local workforce Loss of student enrollment to online schools that specialize and are pure players in distant learning and virtual education delivery

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

A11 Use-case actors

Which stakeholders influence cloud buying decisions? Are there common attributes of these stakeholders and business responsibilities? Here are the key business and IT actors that appear in the use-cases, with their roles, and their impacts on cloud buying decisions.

C-level stakeholders such as the CIO, CTO, and COO rationalize investments in IT for supporting the business processes and improving employee productivity using internal and external collaborations. These stakeholders are constantly seeking new ways to optimize investments, especially when it comes to large-scale investments in end-user services that involve high costs and complexity.

Business Operations Controllers and Business Procurement Managers view cloud as a means to modernize business processes and ease management of business activities; for example, through automatic or self-service IT management facilities. While cloud provides the capability to scale quickly and cheaply, the business operations manager needs to optimize resource utilization and reduce the operating costs. The Business Procurement Manager, on the other hand, is responsible for keeping each project's usage within its budget and reducing end-user costs.

Line of Business (LoB) leaders and IT delivery teams require the assistance of IT to support rapid deployment and launch of new products and services. These teams are required to work collaboratively to manage cloud and integrate this new delivery model in their enterprise.

End users definitely have an impact on buyer decisions though they are often unaware of how their services are implemented.

The actors in the use-cases are analyzed in this section, using a template that has six elements:

- The name of the actor
- The organization to which the actor belongs
- The use-case category (and the organization's industry sector)
- Relationships
- Responsibilities
- Associations to use-cases

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Actor 1: Business Operations Controller

Organization	Cloud Service Consumer
Category	Cross Industry, Financial Services, Government, Pharmaceuticals
Relationships	Shareholders Business Procurement Management IT Operations Management
Responsibilities	Generate management reports for shareholders to report the business performance (Key Performance Indicators (KPI) and Service-Level Agreements (SLA)) to shareholders Provide business process extensions; e.g., BPO Provide governance structure and control to manage disparate activities for project-based resources Provide necessary disaster and recovery back-up processes Ensure that business activities can be carried out during times of additional complexity, such as mergers and acquisitions Provide business mobility services and support Business productivity requirements to improve community and knowledge development in the organization Provide workflow to connect different business processes and applications to improve information quality flow Introduce quality management tracking and reporting of business processes cutting across a number of business areas and applications and database sources
Associations to Use-Cases	Business Activity Reporting Rapid Business Capacity & Scale Operational Efficiency Business Transaction Assurance and Continuity Research Productivity Rapid Business Process Extension

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Actor 2: Business Procurement Manager

Organization	Service Consumer
Category	Cross Industry, Government, Financial Services, Telecom, Energy and Utilities
Relationships	Business Operations Controller IT Operations Management
Responsibilities	Modernize business processes Procure business process extensions; e.g., BPO Provide management (e.g., automatically or self-service) for business activities Provide scale up and down options for resources with associated temporary capacity for business activities Optimize resource utilization Keep each project's usage within its budget Scale capacity quickly and cheaply Reduce end-user costs/charge-back to customers Use commodity services to reduce operating costs Rationalize business assets through alternative sourcing of assets Use external transition services to assist consolidation and rationalization of merger and acquisition activity Manage the cost of business operations and seek alternative ways to offer common services Procure collaboration tools and skills within the business to meet demand Control transaction costs Reduce supply costs through alternative methods of service supply on common services
Associations to Use-Cases	Modernization Rapid Business Capacity and Scale Operational Efficiency Research Productivity Business Utility/Agility Partner Collaboration

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Actor 3: IT Operations Manager

Organization	Service Consumer
Category	Cross Industry, Financial Services, Government
Relationships	Business Procurement Manager Business Operations Manager Data Center Management
Responsibilities	Modernize infrastructure and applications Provide and manage remote user desktop services Infinite compute and storage capacity on-demand Access to temporary computing processing for business simulation and transaction processing Manage scale up and down options for resources (project-based) with associated temporary capacity for business activities Provide access to rapid deployment of compatible and simple end-user services through browser or end-user device Ensure reliable and responsive IT service provision to support product and service launch Provision capacity for storage of information and to support service delivery Provide environment for typical high-performance computing workloads Allow users to submit distributed workloads Bill each project for their use of compute resources Provide elasticity – access to and rapid deployment of end-user services Use shared services to maximize utilization in exploring virtualization and other options for reducing the carbon footprint
Associations to Use-Cases	Modernization Rapid Business Capacity and Scale Operational Efficiency People Productivity/Workplace Collaboration Business Transaction Assurance and Continuity Research Productivity Business Utility/Agility IT Cost Optimization Business Growth Development, Test, and Validation Support On-demand Business Change Validation Green Footprint

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Actor 4: CIO/CTO

Organization	Service Consumer
Category	Cross Industry
Relationships	IT Operations Manager Line of Business (LoB)
Responsibilities	Improve the internal collaboration (e.g., communication methods) for the business Improve the external collaboration (e.g., customer meetings) for the business Rationalize investments in IT Create common utility services to support merger and acquisition activities Provide services and support on-demand to mobile workers Improve the community and knowledge development in the organization Manage large-scale investments in end-user services with high costs and complexity
Associations to Use-Cases	People Productivity Workplace Collaboration IT Cost Optimization End-User Cost Optimization

Actor 5: Development/Test Team Manager

Organization	Service Consumer
Category	Cross Industry
Relationships	IT Operations Manager
Responsibilities	Provision development and test environments at short notice Redeploy development and test environments once tests have been conducted Decrease/maintain Time-to-market for products Provide on-demand test service Provide on-demand business change validation Provide access to development and test services
Associations to Use-Cases	Development and Test Services Business Growth Development, Test, and Validation Support

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Actor 6: Marketing Manager

Organization	Service Consumer
Category	Financial Services
Relationships	New Product Development Business Operations Controller Business Procurement Management IT Operations Management
Responsibilities	Release new products and/or services Ensure rapid development and launch of new product and/or service specification
Associations to Use-Cases	Business Market Launch

Actor 7: Security and Compliance Manager

Organization	Service Consumer
Category	Financial Services
Relationships	Business Operations Controller Data Center Management IT Operations Management
Responsibilities	Ensure private cloud capability is secure and complies with regulations Ensure access, authentication, repudiation facilities are good Ensure audit and eDiscovery compliance Ensure web security authentication is satisfactory Ensure system is securely partitioned and isolated Ensure good disaster recovery capabilities Ensure business continuity
Associations to Use-Cases	Secure Business Assurance Services

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

A12 Use-case benefits

This section presents the benefits of cloud computing that are exemplified by the use-cases, in tabular form. For each use-case, it shows:

- The essential characteristics, service models, and deployment models of cloud computing that contribute to success (the cloud computing "features")
- The advantages that these features bring
- The benefits to the enterprise resulting from these advantages

Use-Case	Feature	Advantage	Benefit
1. Business Activity Reporting	Measured Service	Availability of reports on service use	The enterprise will produce SLA and OLA reports for the internal and external service providers by combining information on service use with location information, and will report its business performance to shareholders by combining the same service information with business information.
2. Modernization	IaaS and SaaS	Ready availability of infrastructure and applications	The company will modernize its infrastructure and applications by obtaining them as cloud services.
	Public Cloud	Ability to use externally provided cloud services without investment or preconditions	The company will not need skills or training.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Use-Case	Feature	Advantage	Benefit
3. Rapid Business Capacity and Scale	Rapid Elasticity	Ability to scale resources to match demand	The company will have access to infinite compute and storage capacity on-demand.
	On-Demand Self-Service Resource Pooling Rapid Elasticity	Resource agility	The company will have access to temporary computing resources for business simulation and transaction processing, and will have the ability to deploy existing business process resources rapidly, through cloud resource agility.
	Private Cloud	Ability to use the cloud model in-house	The company will have the ability to deploy existing business process resources, including applications and infrastructure.
	Rapid Elasticity	Ability to scale resources to match demand	Its ability to extend its business processes will not be compromised by inability to support them with resources.
4. Operational Efficiency	IaaS On-Demand Self-Service	Ready availability of infrastructure Availability of services on-demand Availability by self-service	The company can use IaaS. Projects that need infrastructure can then obtain it on demand, by self-service.
	Measured Service	Availability of reports on service use	Provider reports on service use will facilitate management of the provisioning of the infrastructure and governance of its allocation and use.
5. People Productivity, Workplace Collaboration	Public Cloud SaaS	Ability to use externally provided cloud services without investment or preconditions Ready availability of applications	The company can obtain remote user desktop services, improved communication methods, and better support for customer meetings through cloud services, to improve its internal and external communications. These improvements will be easy to implement, and secure.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Use-Case	Feature	Advantage	Benefit
6. Development and Test Services	PaaS On-Demand Self-Service Resource Pooling Rapid Elasticity	Ready availability of platform Resource agility	The company will introduce development and test services that can be deployed more quickly than at present by using a cloud development/ test platform.
	Rapid Elasticity	Ability to scale resources to match demand	Cost will be minimized through the ability to scale resources to match demand.
7. Business Transaction Assurance and Continuity	On-Demand Self-Service Resource Pooling Rapid Elasticity Public Cloud	Resource agility Ability to use externally provided cloud services without investment or preconditions	The company will ensure that business activities continue effectively during times of additional complexity, such as mergers and acquisitions, by obtaining the additional resources that it needs on-demand from the cloud.
	Resource Pooling Public Cloud	Ability to use resources in different locations Ability to use externally provided cloud services without investment or preconditions	The company will implement the necessary back-up and disaster recovery processes using infrastructure that is off-premise and in different locations.
8. Business Market Launch	PaaS On-Demand Self-Service Resource Pooling Rapid Elasticity	Ready availability of platform Resource agility	The company will speed up the development of applications, and the launch and deployment of the products, by using a cloud development/deployment platform. The platform services will be responsive and reliable.
	Rapid Elasticity	Ability to scale resources to match demand	The company will have increased information storage capacity because the cloud resources can be scaled to match demand.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Use-Case	Feature	Advantage	Benefit
9. Secure Business Assurance Services	Private Cloud	Ability to use the cloud model in-house	The company can implement a secure private cloud platform that will provide partitioning and isolation and ensure compliance with national regulations and with audit and e-Discovery legislation, and can deploy the applications that it needs on this platform.
	Rapid Elasticity	Ability to scale resources to match demand	Use of a cloud platform will mean that the applications are scalable.
	Resource Pooling	Ability to use resources in different locations	The company will implement disaster recovery capabilities to ensure business continuity using infrastructure that is in different locations.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Use-Case	Feature	Advantage	Benefit
10. Research	Resource Pooling	Dynamic assignment of resources to consumers	The facility will be able to allocate compute resources equitably across a mixture of batch and interactive workloads because the cloud services dynamically assign pooled resources.
	Measured Service	Availability of reports on service use	The facility will be able to bill each project for its use of compute resources, using the providers' reports on service use.
	IaaS	Ready availability of infrastructure.	The facility will be able to support the common programming languages of the day by implementing this support on cloud infrastructure.
	On-Demand Self-Service Resource Pooling Rapid Elasticity	Resource agility	The facility will be able to achieve optimal processing performance for typical high-performance computing workloads, and allow users to submit distributed workloads, because of cloud resource agility.
	Rapid Elasticity	Ability to scale resources to match demand	The facility will be able to optimize resource utilization by matching resources to demand.
	Measured Service	Availability of reports on service use	The facility will be able to keep each project's usage within its budget by monitoring provider reports of service use.
	Rapid Elasticity	Ability to scale resources to match demand	The facility will be able to scale capacity quickly and cheaply to match demand.
	On-Demand Self-Service	Ability of users to provision resources.	The facility will have a level of systems management acceptable to the users because the users will be able to provision the resources that they need on-demand.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Use-Case	Feature	Advantage	Benefit
11. Productivity	SaaS Public Cloud	Ready availability of applications Ability to use externally provided cloud services without investment or preconditions	The agency can use collaboration services available on the cloud to achieve better business productivity through improved community working and knowledge development.
	Broad Network Access	Capabilities can be accessed from a variety of client platforms	The agency will have improved support for mobile users because the cloud services can be accessed from mobile 'phones, laptops, and PDAs.
	Rapid Elasticity	Ability to scale resources to match demand	The agency may be able to achieve reduced end-user costs by being able to scale resources to match demand.
	Measured Service	Availability of reports on service use	The agency will be able to charge costs back to its customers, using providers' reports on service use.
12. Business Utility/Agility	PaaS	Ready availability of platform	The agency will be able to create common utility services to support its activities using a cloud platform.
	Rapid Elasticity	Ability to scale resources to match demand	The agency will have IT elasticity because the cloud gives the ability to scale resources to match demand.
	Broad Network Access	Capabilities can be accessed from a variety of client platforms	The agency will give users better access to services because the cloud capabilities can be accessed from a variety of client platforms.
	On-Demand Self-Service Resource Pooling Rapid Elasticity	Resource agility	The agency will be able to achieve rapid deployment of end-user services through resource agility.
	IaaS	Ready availability of infrastructure.	The agency will be able to rationalize its business assets by using cloud infrastructure to replace in-house resources.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Use-Case	Feature	Advantage	Benefit
13. IT Cost Optimization	IaaS, PaaS, and SaaS Public Cloud	Ready availability of infrastructure, platform, and applications Ability to use externally provided cloud services without investment or preconditions	The agency will be able to rationalize its business assets through alternative sourcing of common cloud infrastructure, platform, and software services.
	Rapid Elasticity	Ability to scale resources to match demand	The agency will be able to achieve cost reductions through the ability to scale resources to meet demand.
	Measured Service	Availability of reports on service use.	The agency will have accurate, consumption-based charges for the services through provider reports on service use.
14. End-User Cost Optimization	SaaS	Ready availability of applications	The agency will be able to provide services and support using SaaS applications.
	Public Cloud	Ability to use externally provided cloud services without investment or preconditions	The agency will not need large-scale investments in end-user services.
	On-Demand Self-Service	Ability of users to provision resources	The agency will be able to provide those services and support on-demand, through the ability of users to provision resources.
	Broad Network Access	Capabilities can be accessed from a variety of client platforms.	The agency will be able to provide the services and support to mobile users, because the cloud services can be accessed from mobile phones, laptops, and PDAs. This access will be secure.
	Rapid Elasticity	Ability to scale resources to match demand	The agency will be able to avoid high costs through the ability to scale resources to meet demand.
	On-Demand Self-Service Resource Pooling Rapid Elasticity	Resource agility	The agency will be able to avoid complexity by taking advantage of resource agility to optimize its resource configuration.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Use-Case	Feature	Advantage	Benefit
15. Business Growth, Development, Test, and Validation Support	PaaS	Ready availability of platform	The government unit will be able to create services for business change validation, solution development, and solution testing, using a cloud platform.
	On-Demand Self-Service	Ability to provision services on-demand.	The services will be available and accessible on demand.
16. Service Delivery	Cloud Model	Ability to be a cloud provider	By becoming a cloud provider, the telecommunications company can provide 'network as a service' and 'transport as a service' to online consumer service companies.
	On-Demand Self-Service	Ability to provision services on-demand	The online consumer service companies will have self-service purchasing and configuration capabilities that are usable by companies with no experience in managing network-based services.
17. Partner Colla-boration	SaaS	Ready availability of applications	The telecommunications company can enable information sharing and business activity through a common collaborative toolset provided by applications available as cloud services.
18. Test Assurance Services	PaaS	Ready availability of platform	The company can introduce new test assurance services implemented on a cloud platform.
	On-Demand Self-Service Resource Pooling Rapid Elasticity	Resource agility	The services can be deployed rapidly on demand because of cloud resource agility.
	Rapid Elasticity	Ability to scale resources to match demand	Costs can be kept to a minimum because of the ability to scale resources to match demand.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Use-Case	Feature	Advantage	Benefit
19. Rapid Deployment	IaaS or PaaS	Ready availability of infrastructure or platform	The company can create its means of service delivery and information exchange to enable rapid deployment, using cloud infrastructure or a cloud platform.
	Broad Network Access	Capabilities can be accessed from a variety of client platforms.	The company will be able to deploy the services to mobile users, because the cloud services can be accessed from mobile 'phones, laptops, and PDAs. This access will be secure.
	On-Demand Self-Service Measured Service	Ability to provision services on demand Availability of reports on service use	The means will be simple to manage because of the ability to provision services on-demand and the availability of provider reports on service use.
20. Rapid Business Process Extension	SaaS Public Cloud	Ready availability of applications Ability to use externally provided cloud services without investment or preconditions	The company may be able to use SaaS products to support the process extension available on the open market, either stand-alone or as an extension to its existing applications.
	IaaS or PaaS	Ready availability of infrastructure or platform	Otherwise, the company can set up its project administration, development environment and test environment, and its workflow engine, using cloud infrastructure or a cloud platform.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Use-Case	Feature	Advantage	Benefit
21. Brand Unification	SaaS	Ready availability of applications	The company can transform its workforce into a globally integrated team by using cloud-based collaboration services. These services can be reliable and secure.
	Private Cloud	Ability to use the cloud model in-house.	The company can leverage existing IT infrastructure without burdening IT resources by implementing a private cloud.
	SaaS Public Cloud	Ready availability of applications Ability to use externally provided cloud services without investment or preconditions	The company may find solutions that may have better capabilities offered as SaaS services.
	PaaS	Ready availability of platform	Alternatively, by incorporating PaaS in its private cloud, the company will enable software providers to port their products to its environment very easily.
	On-Demand Self-Service Resource Pooling Rapid Elasticity	Resource agility	Resource agility will contribute to timely order fulfillment.
22. Innovation – New Ventures	IaaS On-Demand Self-Service Resource Pooling Rapid Elasticity Public Cloud	Ready availability of infrastructure Resource agility Ability to use externally provided cloud services without investment or preconditions	The company can obtain hardware, software, and operating systems through IaaS; resource agility will then enable it to implement more efficient processes for purchasing, configuring, and ongoing maintenance, and achieve provisioning and de-provisioning of servers in minutes.
	SaaS Public Cloud	Ready availability of applications Ability to use externally provided cloud services without investment or preconditions	The company may find the applications that it needs to deliver improved business services offered as SaaS services.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Use-Case	Feature	Advantage	Benefit
23. Green Footprint	Rapid Elasticity	Ability to scale resources to match demand	The companies can maximize utilization by scaling resources to meet demand. (The ability to do this is based on virtualization and shared services.)
	Resource Pooling	Ability to use resources in different locations	The companies can reduce their carbon footprint (or, at least, appear to do so) by using off-premise resources.
24. Virtual Learning	Community Cloud	Ability for a community to share infrastructure	The colleges can share their virtual learning applications and infrastructure by implementing a community cloud.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Appendix B

Glossary

Added Service	A service performed using cloud 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.
ADM	Architecture Development Method; see [TOGAF].
AMPU	Average Margin Per User. The margin made by the provider from each customer, typically measured as the revenue minus the costs and divided by the number of users. This is a widely-used metric for gauging the success of businesses in the telecommunications industry.
API	Application Programming Interface.
Appliance	A self-contained IT system that can be plugged into an existing IT infrastructure to carry out a single purpose. An appliance is designed to address a specific IT operation from within a closed architecture that may contain an operating environment, storage, and specific applications. The appliance's purpose could be to provide additional processing power, network storage or monitoring, or anti-virus and security.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

ARPU	Average Revenue Per Unit. A measure of the revenue generated per user or unit. Average revenue per unit allows for the analysis of a company's revenue generation and growth at the per-unit level, which can help investors to identify which products are high or low revenue-generators. This measure is most often used in the telecommunications sector to survey the amount of revenue generated per cell-phone user, for example.
ASP	Active Server Page. A kind of dynamically-generated web page in Microsoft's .NET framework.
Authentication	The process of establishing confidence in the truth of some claim. In the context of identity management, an authentication system provides an understood level of confidence that an identifier refers to a specific individual (individual authentication) or identity (identity authentication), or that an attribute applies to a specific individual (attribute authentication).
Availability	The proportion of the time that a system is available for use. It is typically measured in 9s. A 'Five 9s' system is up 99.999% of the time – a little over five minutes per year downtime. Planned, scheduled outages for maintenance are typically excluded.
BASEL II	The second set of recommendations on banking laws and regulations issued by the Basel Committee on Banking Supervision. It deals with how much capital banks need to put aside to guard against financial and operational risks.
BI	Business Intelligence.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

BPaaS	Business Process as a Service. A service model may be added to the original models of IaaS, PaaS, and SaaS, in which the consumer has the ability to use provider-defined business processes running on a cloud infrastructure. The business processes interface with various client devices through a thin client interface such as a web browser. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual business processes and underlying application capabilities, with the possible exception of limited user-specific process configuration settings.
BPM	Business Process Management.
BPO	Business Process Outsourcing.
Broad Network Access	The essential characteristic of cloud computing that 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 PDAs).
Business Operations Controller	A person that controls or manages ongoing recurring activities involved in the running of a business for the purpose of producing value for the stakeholders. Refer to Business Operations Controller (in Appendix A) for a description of the typical responsibilities of a Business Operations Controller.
Business Procurement Manager	A person that manages the sourcing and purchasing of goods and services for business use. Refer to Business Procurement Manager (in Appendix A) for a description of the typical responsibilities of a Business Procurement Manager.
Buyer	A person or organization that contracts to acquire an asset in return for some form of consideration. For cloud computing, the contracted asset will be a cloud computing service; e.g., a virtual machine for a period of time or use of a cloud computing software application.
CAPEX	Capital Expenditure.
CFAT	Cash Flow after Taxes.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

CEM	Customer Experience Management.
CEO	Chief Executive Officer.
CFO	Chief Financial Officer.
CiA	Capacity in Advance. An arrangement under which a service provider maintains spare capacity at its cost and only charges for what is used.
CIO	Chief Information Officer.
Cloud Burst	A technique used by hybrid clouds to provide additional resources to private clouds on an as-needed basis. If the private cloud has the processing power to handle its workloads, the hybrid cloud is not used. When workloads exceed the private cloud's capacity, the hybrid cloud automatically allocates additional resources to the private cloud.
Cloud Computing	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.
Cloud Service	A service offered over a network or the Internet that has the essential characteristics of cloud computing.
CMU	Carnegie-Mellon University.
CoD	Capacity on Demand.
Community Cloud	The cloud computing deployment model in which 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.
Compliance	The act of adhering to, and demonstrating adherence to, a standard or regulation.
Consumer	A person or organization with needs who makes use of a service.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

COO	Chief Operations Officer.
CoS	Cost of Service.
COSO	Committee of Sponsoring Organizations of the Treadway Commission. See [COSO].
CPU	Central Processing Unit.
CRM	Customer Relationship Management.
CSA	Cloud Security Alliance. See [CSA].
CTO	Chief Technical Officer.
DEV	Development.
DIACAP	The US Department of Defense Information Assurance Certification and Accreditation Program. It defines the risk management process for DoD information systems.
DMTF	Distributed Management Task Force. See [DMTF].
DoD	The Department of Defense of the USA.
Driver (risk)	A factor that has a strong influence on the eventual outcome or result.
EC	European Commission.
Eclipse	A particular open source software development environment. See [ECLIPSE].
Ecosystem	A system of participants in a defined market that have integrated business processes and use common standards for exchange of information, products, and services.
e-Discovery Legislation	Legislation relating to discovery in civil litigation which deals with the exchange of information in electronic format.
EOL	End of Life.
End User	A person who uses a service and interacts with its performer. End users are often unaware of how their services are provided or procured. Examples of these actors include employees, mobile users, and World-Wide Web users.
Enterprise	A collection of organizations that has a common set of goals and/or a single bottom line. An enterprise can be a government agency, a whole corporation, a division of a corporation, a single department, or a chain of geographically distant organizations linked together by common ownership.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

ERM	Enterprise Risk Management.
ESB	Enterprise Services Bus.
EU	European Union.
Factor (risk)	A component of a system that significantly affects its behavior.
FedRAMP	Federal Risk and Authorization Management Program (in the USA). See [FEDRAMP].
FISMA	The US Federal Information Security Management Act of 2002.
FM	Financial Management.
Fremium	A business model in which a provider gives a service away free of charge, possibly but not necessarily supported by advertising, acquires a lot of customers very efficiently through word of mouth, referral networks, organic search marketing, etc., then offers premium priced value added services or an enhanced version of the service to the customer base.
GLBA	The US Gramm-Leach-Bliley Act. This Act is also known as the Financial Services Modernization Act of 1999. It provides limited privacy protections against the sale of private financial information.
Grid Computing	A computing model in which a wide variety of geographically distributed computational resources are presented as a single, unified resource. A computer grid is analogous to an electric power network (grid), where power generators are distributed, and consumers are able to use electricity without bothering about where it comes from. Grid computing can be used to solve problems that require computationally intensive processing, as in the SETI@ home project.
HIPAA	The US Health Insurance Portability and Accountability Act (HIPAA) of 1996. It includes provisions that address the security and privacy of health data.
HPTC	High Performance Technical Computing.
HW	Hardware.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Hybrid Cloud	The cloud computing deployment model in which 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.
Hypervisor	A particular virtualization technique that allows multiple operating systems, termed guests, to run concurrently on a host computer.
IaaS	Infrastructure as a Service. The cloud service model in which 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).
ICAEW	Institute of Chartered Accountants in England and Wales. See [ICAEW].
IEC	International Electrotechnical Commission.
I/O	Input/Output.
IP	Internet Protocol.
ISAE 3402	The International Standards for Assurance Engagements (ISAE) No. 3402. This standard applies to the assessment and audit of the internal controls of service organizations.
ISO	International Organization for Standardization.
IT	Information Technology.
ITIL	Information Technology Infrastructure Library. See [ITIL].
JSP	Java Server Pages.
K	Kilo.
KPI	Key Performance Indicator. Indicator to measure the defined business operations goals and provide opportunities to further optimize the business to achieve business objectives.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

LoB	Line of Business.
M	Mega.
Measured Service	The essential characteristic of cloud computing that 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.
Metric	An attribute that can be measured. Metrics are used to measure the QoS performance and the state of progress of people and organizations.
Metering	A mechanism to measure usage of resources in cloud computing. It is usually measured per client and application with a measurement frequency of daily, weekly, monthly, and annually. Metering is also used by the cloud consumer in comparative study of cloud vendors on cost and reliability.
Mission Risk	A systemic risk that affects a program's ability to achieve its key objectives.
MIT	Massachusetts Institute of Technology.
Monitoring	A mechanism to collect information on services performances in a controlled environment for further analysis.
ms	Milliseconds.
MTBF	Mean Time Between Failures. Used in conjunction with MTTR as a measure of reliability for repairable components and, by analogy, for cloud services.
MTTF	Mean Time To Fail. A measure of reliability used for non-repairable components.
MTTR	Mean Time To Repair. Used in conjunction with MTBF as a measure of reliability for repairable components and, by analogy, for cloud services.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Multi-Tenancy	A principle in software architecture where a single instance of a computing resource serves multiple client organizations (the tenants) providing a separate environment for each. How cloud 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.
NAS	Network-Attached Storage. File-level computer data storage connected to a computer network providing data access to heterogeneous clients.
NIST	National Institute of Standards and Technology (of the USA). See [NIST].
NOC	Network Operations Center.
NPV	Net Present Value.
OAuth	Open Authorization. See [OAuth].
Object-orientation	Often used in connection with object-oriented programming, a programming paradigm using software constructs called 'objects'.
OGF	Open Grid Forum. See [OGF].
OLA	Operational-Level Agreement.
On-Demand Self-Service	The essential characteristic of cloud computing that 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.
OPEX	Operational Expenditure.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

OS	Operating System. A collection of programs and data that efficiently manages computing resources.
Outsourcing	A service model in which the consumer transfers the management or execution of the entire IT function to an external provider for a regular (e.g., monthly) fee.
Outtasking	A service model in which the consumer turns over a specific IT operation to the provider organization, rather than the entire IT function.
PaaS	Platform as a Service. The cloud service model in which 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.
Pay-as-you-go	A financial policy by which capital projects are financed from current revenue in the operating budget rather than through borrowing. For cloud computing, an arrangement in which payment for IT resources is on a per-use-time basis using an OPEX style charging principle.
Pay-by-the-drink	Usage-based billing.
Pay-upfront	A billing model in which payment is made in advance.
PC	Personal Computer.
PCI DSS	The Payment Card Industry Data Security Standard. It states international security requirements for protecting cardholder data. See [PCI-DSS].
PDA	Personal Digital Assistant.
Peer-to-Peer Network	A computer network model in which each node plays a similar role. Peer-to-peer networking is contrasted with other models where different nodes play different roles. For example, in the client-server model, clients and servers are not interchangeable.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

PHP	Hypertext Preprocessor. PHP is a widely-used general-purpose scripting language that is especially suited for web development.
PLM	Product Lifecycle Management.
Pricing Model	A financial model used to determine pricing of a product or service.
Private Cloud	The cloud computing deployment model in which 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.
Process	A sequence or flow of activities in an organization with the objective of carrying out work.
Provider	A person or organization that offers a capability exposed as a service. A cloud computing provider is a person or organization that offers capabilities exposed as a service that satisfies the essential characteristics of cloud computing. 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.
Public Cloud	The cloud computing deployment model in which the cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.
QoS	Quality of Service. Outlines the non-functional requirements (e.g., performance) of a resource in a controlled computing environment.
R&D	Research and Development.
Rapid Elasticity	The essential characteristic of cloud computing that 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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Resource Pooling	The essential characteristic of cloud computing that 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.
Risk	The effect of uncertainty on objectives.
Risk Management	The activities and methods used by an organization to manage risks associated with its business operations.
ROI	Return On Investment.
RPO	Recovery Point Objective.
RTO	Recovery Time Objective.
SaaS	Software as a Service. The cloud service model in which 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.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Safe Harbor	A US government program that helps US companies to establish adequate compliance with data privacy laws in other countries. The European Commission's Directive on Data Protection went into effect in October 1998, and would prohibit the transfer of personal data to non-European Union countries that do not meet the European Union (EU) 'adequacy' standard for privacy protection. The Safe Harbor program provides a streamlined means for US organizations to comply with this Directive. It also provides similar help with regard to data privacy legislation in Switzerland.
Sarbanes-Oxley	A US federal law enacted on July 30, 2002, which sets standards for all US public company boards, management, and public accounting firms, in particular for corporate accounting.
SAS 70	Statement on Auditing Standard 70. A Report on the Processing of Transactions by Service Organizations where professional standards are set up for a service auditor that audits and assesses internal controls of a service organization. It was developed and is maintained by the American Institute of Certified Public Accountants.
SCM	Supply Chain Management.
Scorecard	A simple record in which results can be entered. Paper scorecards are used in games such as golf to keep score. The term is used in IT jargon to refer to a record, which may be held electronically, that describes some aspect of a person, organization, system, or activity. A Balanced Scorecard (BSC) is a particular form of scorecard: a semi-standard structured report supported by proven design methods and automation tools that can be used by managers to keep track of the execution of activities by staff within their control and monitor the consequences arising from these actions.
SEI	Software Engineering Institute (of the Carnegie-Mellon University).

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Service	A logical representation of a repeatable business activity.
Service Contract	The contract between the provider and the consumer of a service, relating to its provision and consumption. 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.
Service-Oriented Architecture (SOA)	An architectural style that supports service-orientation, which is a way of thinking in terms of services and service-based development and the outcomes of services.
SETI@home	(SETI at home) an Internet-based public volunteer computing project to search for extra-terrestrial intelligence, using the volunteers' computers as a distributed computing grid to analyze radio signals.
SFA	Sales Force Automation.
Single Sign-On (SSO)	An access control method in which multiple systems with individual access control mechanisms are federated in such a way that a user can sign on to one of them and then access them all without signing on again.
SLA	Service-Level Agreement. SLAs are characteristics of services that are met by service providers and are expected by service consumers as per defined contractual agreements.
SME	Small and Medium-sized Enterprises.
SMP	Symmetric Multi-Processing.
SNIA	Storage Networking Industry Association. See [SNIA].
SOA	Service-Oriented Architecture.
SOE	Standard Operating Environment. A standard operating system configuration.
SOX	Sarbanes-Oxley.
SPOF	Single Point of Failure.
SPL	Software Product Line.
SQL	Structured Query Language.
SRM	Supplier Relationship Management.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

SSAE 16	Statement on Standards for Attestation Engagements (SSAE) No. 16. This standard applies to the assessment and audit of the internal controls of service organizations. It mirrors and complies with ISAE 3402 and effectively replaces the Statement on Auditing Standard 70 (SAS 70) Type II.
SSL	Secure Sockets Layer. A secure data communications protocol.
SSO	Single Sign-On.
TCO	Total Cost of Ownership.
TCP/IP	Transmission Control Protocol/Internet Protocol.
TOGAF	The Open Group Architecture Framework. See [TOGAF].
Throughput	The amount of work that a computer can do in a given time period.
UI	User Interface.
UNIX	The UNIX System. See [UNIX].
US	United States (of America).
USA	United States of America.
Utility Computing	The packaging of computing resources as a metered service similar to a traditional public utility.
Virtual Machine	A software program that emulates a physical machine.
Virtualization	A means of separating the execution of a software environment from the underlying hardware. Virtualization may be: *hardware-assisted* (where, as the term suggests, the hardware has specific features within its design to aid virtualization), *full-virtualization* sufficiently separating software execution from the hardware resources so as to allow operating systems to operate unmodified, *partial-virtualization* where some aspects of software execution require modification to run successfully, or *para-virtualization* where for performance reasons a specific port of an operating system or suitable virtual machine interface is required for communicating with hardware in a virtualized environment to offset the performance degradation associated with executing certain functions in a virtualized machine *versus* on the native hardware.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

VPN	Virtual Private Network.
Web 2.0	The second generation of web sites, having greater sophistication than simple web pages. The greater sophistication takes several different forms; there is no agreed precise definition of what characterizes a Web 2.0 site.
XML	Extensible Markup Language.
XMP (Cray)	A model of supercomputer produced by Cray Research.

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Index

A

access control
  end user 94
  provider 95
added services 5, 22
ADM 195
agility 35
AMPU 195
API 148
appliance 195
architecture vision 47
ARPU 196
ASP 196
asset management 137
asset sharing 3
authentication 27
availability 89, 122

B

back-up 87
BASEL II 196
BNA 38
BPaaS 18, 139
BPM 30
BPO 57
broad network access 5
bulk data transfer 88
business context 48
business continuity 41
business continuity management 118
business intelligence 68
Business Operations Controller 178
business opportunities 45
business paradigm 28
business process 18, 35
Business Procurement Manager 178
buyer 4

C

CAPEX 42, 44, 62, 80
carbon footprint 41
CEM 198
CEO XVI, 126
CFAT 197
CFO XVI
CiA 198
CIO XVI, 178
cloud burst 17
cloud buying lifecycle 75
cloud computing XIII, 1
  definition 2
  deployment models XV
  essential characteristics XV, 3, 4
  model 3
  resources 3
  service models XV
cloud ecosystem 30
cloud provider 139
cloud resources 1
cloud roadmap 48

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

cloud service 4, 5, 14, 18, 20, 25, 29, 38, 75, 198
cloud service provision 45
cloud vision 47
CMU SEI XX, 109
CoD 106
cold start 92
collaboration 38, 54
community cloud 16, 44
competitive pressure 140
compliance 118
configurability 93, 124
consumer 4
contract 100
COO 178
CoS 199
COSO 110
cost 42
cost model 77, 82, 136
cost sharing 44
CPU 94, 147
Cray 210
CRM 18, 20, 157
CSA 124
CTO XVI, 178

D

decision tree 57
deployment 135
deployment models 4, 15
DEV 199
DIACAP 97
DMTF 89
DoD 97
driver 111

E

EC 55
Eclipse 24
ecosystem 110, 139, 156
eDiscovery legislation 199
end user 17, 91, 94, 178
enterprise XIII, 1, 33
EOL 199
ERM 110
ESB 65
EU 97, 98, 126
Eucalyptus
  open source platform 13
exit strategy 98
external service 120

F

factor 26
FAIR 124
fault management 94, 124
fault tolerance 122
FedRAMP 109
financial risk 115
FISMA 97
Fremium 200
functionality 122

G

GLBA 97
grid computing 3, 200

H

HIPAA 97
HPTC 200
HW 200
hybrid cloud 17, 43, 57
hypervisor 13

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

I

IaaS XV, 3, 4, 12, 37, 67, 76
ICAEW 110
IEC 201
information restriction 54
information silo 27
Internet 12
I/O 77
ISAE 97
ISO 31000:2009 110
ISO/IEC 27005 124
ITIL 85, 96
IT provision 40

J

Java 14, 89
JSP 201

L

legislation 97
linear scaling 9
Linux 14, 86, 89
LoB 178
logging 96
long tail 139

M

manageability 40, 124
measured service 10
metering 5
metric 24, 147
mission risk 110, 113
MIT 13
modeling resources 81
monitoring 103
Mosaic 109, 110
MTBF 90, 122
MTTF 90
MTTR 90, 122
multi-tenancy 7, 93, 133

N

NAS 3
negotiation 101
NIST 1, 2
NIST model 18
NOC 203
NPV 203

O

O-ACEML 98
OAuth 95
object-orientation 138
OGF 89
OLA 203
on-demand self-service 4
on-demand utilization 9
OPEX 9, 42, 44, 62
organizational culture 115
outsourcing 12, 37, 57
outtasking 204
ownership 80

P

PaaS XV, 3, 13, 37, 67, 76
pay-as-you-go 9, 80, 134
pay-by-the-drink 148
pay-upfront 45
PCI DSS 97
PDA 6
peer-to-peer network 2
performance 122
PHP 14
PLM 205
pricing model 24
private cloud 15

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

process 28
productivity 38, 132, 147
provider 3
public cloud 16

Q

QoS 6, 33, 52, 62, 140
quality 39, 141, 148

R

rapid elasticity 9
R&D 205
recoverability 90, 122
regulations 97
reliability 90
reporting 94, 124
requirements 85, 105
resource partitioning 96
resource pooling 6
resource scaling 36
resource utilization 133
response time 123
responsiveness 91
risk 109
risk assessment 110, 125
risk communication 111
risk management 110
risk mitigation 127
risk-reward balance 48
ROI XX, 131, 132, 141
RPO 62, 90
RTO 62, 90

S

SaaS XV, 3, 14, 37
Safe Harbor 207
Sarbanes-Oxley 97
SAS 70 97, 207
scale 135, 148
SCM 64
scorecard 106, 149
security 124
server consolidation 42
service
  procurement 128
  selection 128
service contract 76
service functionality 86
service integration 116
service models 3, 11
SETI@home 3
SFA 208
shared logic 38
single sign-on 27
SLA 11, 22, 62, 90, 121
SME 110, 208
SMP 208
SNIA 89
SOA 2, 30, 65, 138
SOE 208
solution architecture 126
solution operation 129
SOX 97
SOX II 52
specialization 135
speed of operation 135, 148
SPL 208
SPOF 122
SQL 14, 208
SRM 208
SSAE 16 97
SSL 209
SSO 209
stakeholder 178

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

standards 6, 30
supplier choice 88
system quality 119, 122

T

TCO 52, 62
TCP/IP 88, 209
The Open Group XXII
thin clients 44
threat management 96
throughput 77, 93, 123
time compression 148
TOGAF 48, 112, 126

U

UNIX 86, 89
usage-based pricing 133
usage information 39
use-case 33, 159
  actors 178
  benefits 184
user interface 44
user satisfaction 125
utility computing 141, 209
utilization 147

V

virtualization 7, 12, 39, 57, 80
virtual machine 4
VPN 210

W

Web 2.0 30, 138
Windows 14
workload 103
  allocations 78
  factors 77
  model 77

X

XML 94
XMP 210

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net

Copyright protected. Use is for members of The Open Group only.
For information and printed versions please see www.vanharen.net