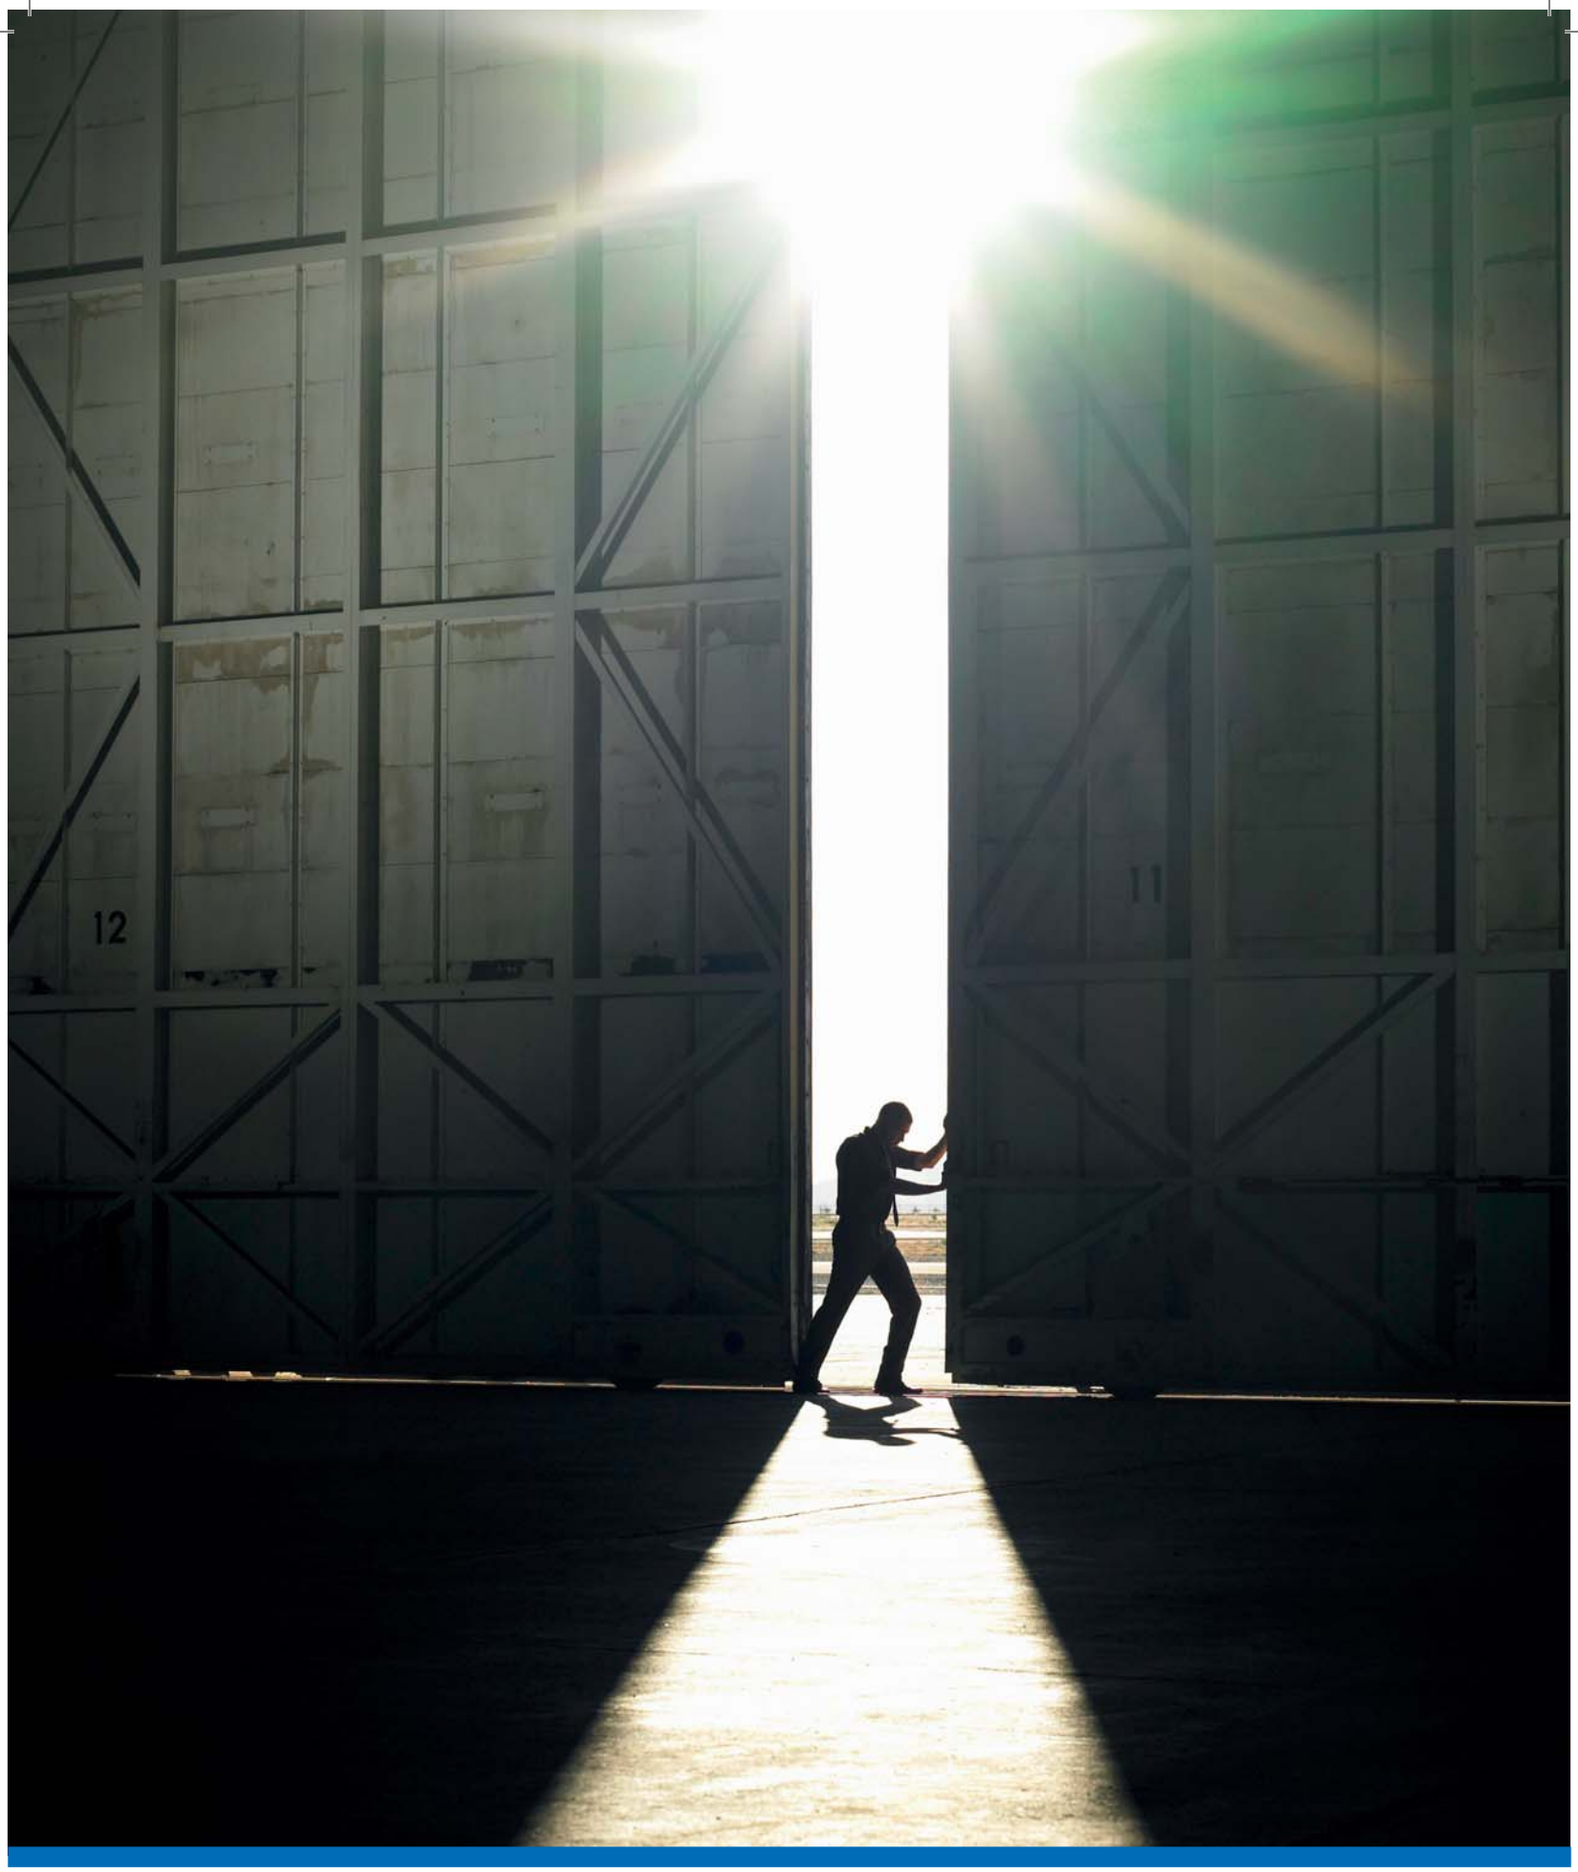

Open for business

The value of Open Source Software in transaction processing

Professor Jim Norton FREng
September 2012

Sponsored by:

amadeus
Your technology partner

Contents

About the author

Professor Jim Norton FREng BEng CDir CITP CEng DEng LLD FIoD FBCS FIET FRSA is the Immediate Past President of the BCS – the Chartered Institute for IT, a board member of the Parliamentary Office of Science and Technology, Chief External Examiner for the Institute of Directors' Certificate in Company Direction and a trustee of the Foundation for Information Policy Research.

His past roles include Chief Executive of the UK Radiocommunications Agency and Chairman of Deutsche Telekom Ltd. He is a respected adviser to governments, an independent director and has served on several notable enquiries and commissions. He led the team that authored the report on e-commerce for the British Prime Minister and was a member of the IPPR's Commission on National Security in the 21st Century.

Any views expressed in this paper are the author's own and do not necessarily represent the opinion of Amadeus IT Group.

Foreword

The shift to Open Source Software is one of the most important current trends in technology, yet it is surprisingly little discussed compared with other developments – including those which, ironically, have only been made possible by the adoption of open systems.

Open source can offer huge benefits, enabling faster innovation and reduced total cost of ownership. Whilst transitioning from closed to open systems is no trivial task, unless this step is taken, businesses risk being left behind as their competitors capitalise on the new possibilities this offers.

Amadeus IT Group is committed both to open source and to thought leadership. We regularly commission white papers on subjects we consider important to the future development of our industry and are delighted to sponsor this report by the distinguished technology leader, Professor Jim Norton. We hope it will both prove informative and also promote debate on this vital, yet sometimes overlooked, issue.

Hervé Couturier
Executive Vice President, Development
Amadeus IT Group

Amadeus IT Group and open systems

Amadeus IT Group is a leading transaction processor and provider of advanced technology solutions for the global travel and tourism industry. Innovation is key to the company's success – the European Commission lists it as the number one investor in research and development (R&D) for the computer services and travel & tourism sectors, with over €340m invested in R&D in 2011.

Amadeus IT Group is strongly committed to open systems at several levels.

- At the operating system level, an ongoing transition from proprietary systems to Linux now sees Altéa Reservations, Inventory and Departure Control System substantially running on open systems.
- Many leading innovations including e-Retail, Dynamic Website Manager and Extreme Search have been built with extensive use of open source components.
- Amadeus ARIA™ Templates – the framework upon which Amadeus IT Group's web solutions are based – has recently been made open source, allowing third party developers to use it without charge.

1. Executive summary and key findings

1.1 This white paper charts the development of Open Source Software (OSS) and open source operating systems from their beginnings 20 years ago. At that time they were seen as upstarts, unlikely to offer a serious challenge to the established orthodoxy of proprietary systems. Now they are most definitely mainstream, adopted and championed by many of the IT industry leaders who initially discounted them.

1.2 The paper asserts that OSS is now applicable in even the most demanding real-time processing environments – such as that seen in the global travel industry.

1.3 Drawing on the experience of Amadeus IT Group, and others with highly demanding applications, the benefits of OSS are explored, including for the business customer:

- greater innovation;
- faster response to change; and
- the ability to support a wide range of heterogeneous systems.

1.4 Similarly the benefits to business providers include:

- better access to skilled, motivated and innovative development and support staff;
- faster exploitation of new technology developments;
- the ability to draw on a global community for specialist knowledge and problem solving;
- avoiding dependency on monopoly suppliers;
- reduced total cost of ownership; and
- full visibility of, and thus confidence in the source code used.

1.5 The paper recognises that the transition from closed to open systems represents a major challenge. It summarises lessons learnt including:

- the importance of maintaining common development/support teams between existing and replacement systems;
- recognising the extent of the challenge and the substantial human and financial resources required;
- the importance of being realistic about the substantial period of time required;
- the continuing challenge of ensuring data integrity during failure recovery or system rollback; and
- the need for advanced filtering and processing of warnings and alarms to ensure that it is possible to “see the wood for the trees”.

1.6 The challenges for auditors and compliance officers in verifying integrity and fitness for purpose in distributed systems is also explored. Where key functions can move between countries or even continents on demand, new approaches and working relationships are required.

1.7 Finally, the paper examines the future potential of OSS and cloud computing. It asserts still further impact across a broad swathe of computing and informatics. In particular the paper explores the:

- work in progress to improve the efficiency of both public and private cloud platforms in demanding applications.
- opportunities for “hybrid” operation merging both private and public clouds;
- challenge of interoperability and portability between public cloud providers; and
- trend towards applications that are “neutral” and not dependant on specific operating systems or browsers.

2. Introduction

2.1 A brief chronology of open systems

2.1.1 Open systems in computing and informatics refers to a class of systems which are built using Open Source Software (OSS) standards and that offer a good level of portability and independence from the hardware platforms on which they operate. They usually include a right to edit and redistribute and have particular characteristics including:

- no intentional secrets: the standard must not withhold any detail necessary for interoperable implementation.
- availability: the standard must be freely and publicly available under royalty-free terms at reasonable and non-discriminatory cost.
- patents: all patents essential to implementation of the standard must be:
- licensed under royalty-free terms for unrestricted use, or
- covered by a promise of non-assertion when practiced by open source software.

2.1.2 There are three key types of licence under which OSS may be released:

- the GNU[i] General Public License (GPL) requires that altered or extra code added to GPL software be also licensed under the GPL. This ensures the propagation of OSS but can cause licensing conflicts if GPL and proprietary software are combined.
- the Berkeley Software Distribution (BSD) licence gives anyone the freedom to release updates or modifications of the software under any licence they wish.
- the Lesser GPL (LGPL) is a compromise between the restrictive GPL and the permissive BSD. Altered LGPL software must continue under LGPL, but extra code can be added under almost any licence the author wishes.

2.1.3 The concept of open systems dates back to the late 1960s and early 1970s, as the first steps were taken to link heterogeneous computer systems together across communications networks. Networking developments took place principally in two communities, the defence/academic world, leading to the development of ARPANET and ultimately to today's Internet, based on the TCP/IP[ii] protocols, and by the public telecommunications operators using a different protocol stack (the "X" series) to achieve similar aims. The reality of multiple competing solutions led to the creation of a single Open Systems Interconnection (OSI) reference model.

2.1.4 Once open solutions were available for networking, attention turned to the attached computer systems. Could the software, both operating system and applications, be made more independent of the particular hardware platform or vendor used? The most widely recognised precursor of today's open operating systems came again from the telecommunications world. AT&T had first developed "Unix" in 1969 at its Bell Laboratories subsidiary. In 1973 Unix was very largely re-written in the high level programming language "C" making it much easier to transfer between multiple hardware platforms. In its later incarnations in the 1980s, Unix embodied standardised programming interfaces and management of attached devices.

2.1.5 Unix was the inspiration for the development of Linux the basis for the vast community of open software and applications in use today. The first Linux Kernel (operating system core) was released by Linus Torvalds in 1991. Some 20 years of development and operating experience, sourced from around the globe, have now gone into the Linux based open operating systems.

2.2 From closed to open systems: a spectrum

2.2.1 It would be wrong to regard the question of whether an operating system is open or closed as simply a binary choice. The reality is that most practical implementations lie somewhere on the spectrum between these two extremes. Many current systems draw on the Unix/Linux legacy whilst arguably now being closed. Examples include the Apple operating systems OSX and iOS,

i GNU is an open source operating system

ii Transmission Control Protocol/Internet Protocol

Oracle Solaris and Blackberry BBX. Others build in their own proprietary additions for example IBM's AIX and HP's UX. Still others remain very largely open such as Google's Android.

2.2.2 Even where a core Linux system is used, there will often be a need to purchase maintenance and support services. The major vendors, such as Red Hat and SUSE, have every incentive to build in some competitive differentiation by customising their various services and tools, particularly in the area of system management.

2.3 Virtualisation: from open systems to cloud computing

2.3.1 Virtualisation is again a development first introduced in the 1970s. A single mainframe computer could run many separate instances of the same operating system at the same time under the control of a "hypervisor"; in effect each instance of the operating system simulated a stand-alone machine. All were hosted simultaneously, but separately, on the same hardware. The use of such virtualisation techniques with Linux style open systems has generated major benefits with single machines supporting multiple virtual servers. For applications that have relatively low load factors the advantages are substantial, lower capital costs through better utilisation, plus ease of management and provisioning.

2.3.2 The final step to create cloud computing was the development of the technique known as "orchestration". This allows many virtual machines to be managed across a range of physical hardware. As load increases for a particular application, further virtual machines can be started automatically and, if necessary further physical systems assigned to support of these in-demand virtual machines and applications.

2.3.3 Almost all of the successful web businesses launched in the last 10 years have drawn on this base of virtualised open systems: familiar names such as Google, Amazon, YouTube, Facebook and Twitter all operate in this way.

> **"The meaning of open".** In a post to the official Google Blog on 21st December 2009, Google Senior Vice President for Product Management Jonathan Rosenberg wrote: "... At Google we believe that open systems win. They lead to more innovation, value, and freedom of choice for consumers, and a vibrant, profitable, and competitive ecosystem for businesses..." and "...Complacency is the hallmark of any closed system. If you don't have to work that hard to keep your customers, you won't... "

2.3.4 The combination of virtualised open systems plus sophisticated orchestration leads to the concept of purchasing computing capability as a utility – so-called cloud computing.

2.4 Stages of the cloud

2.4.1 The terminology of cloud computing is littered with confusing four-letter acronyms. Essentially, there are four stages in cloud provision. At each successive stage an increasing level of functionality is sourced from the cloud service provider. The four stages are well described in the Institution of Engineering and Technology's Factfile, Cloud computing as:

- Infrastructure as a Service (IaaS). At this stage basic virtual machines, virtualised storage and network access are made available to users. Those users are responsible for providing and loading their own operating systems, management systems, middleware and applications. Commercial examples include S3 from Amazon, BlueCloud from IBM, Project Caroline from Sun and similar offerings from both Google and Microsoft.

- Platform as a Service (PaaS). This adds to IaaS a pre-configured operating system and middleware. The user remains responsible for providing and managing their own applications. Commercial examples include EC2 from Amazon, App Engine from Google, Windows Live from Microsoft, Force.com from Salesforce.com and SuiteFlex from NetSuite.

- Software as a Service (SaaS). This builds on IaaS by transferring responsibility for the application software to the cloud service provider. Commercial examples include Office Live from Microsoft, Zimbra from Yahoo Salesforce.com and Google Apps.

- Business Process as a Service (BPaaS). At this final point the cloud service provider would provide a complete suite of applications to perform a major business function or

functions, including taking responsibility for integration between those applications. Commercial examples include Infosys BPaaS and IBM SmartCloud.

Growth of the cloud: The latest research from Gartner suggests that "cloud computing" is the fastest growing service type in the outsourcing market achieving 48% growth last year with total spending of $5bn. Gregor Petri, research director at Gartner noted that: "As next-generation business applications come to market and existing applications are migrated to use automated operations and monitoring, increased value in terms of service consistency, agility and personnel reduction will be delivered [through the cloud]"

2.4.2 Cloud can be provided entirely in-house for example by an existing corporate IT Department – the so called private cloud. It can also be offered under contract by an external service provider such as Amazon or Google – the so-called public cloud. Figure One below illustrates these four stages in both private and public cloud provision. Other combinations and hybrids between private and public provision are possible but the basic elements are as outlined above.

Figure One: an illustration of the key stages of both private and public cloud provision

The areas shaded in green indicate the elements transferred to provision by the cloud operator under the various different scenarios.

With grateful acknowledgement to the Institution of Engineering and Technology.

Quick take:

Open systems are now a strong and proven approach:

- OSS has a 40-year pedigree, with Linux launched more than 20 years ago.
- The distinction between open and closed systems is not absolute. Existing major closed systems, such as Apple's OSX and iOS, draw on open roots such as UNIX.
- Virtualised open systems represent a key stepping stone to private and public computing on demand through cloud technologies.
- Cloud computing itself can comprise several stages, with progressively more elements being contracted out to the service provider.

3. Open systems in transaction processing

3.1 As the use of open operating systems and applications has matured, they have increasingly been applied in business critical applications. The open source Apache Web server has been the most widely used such server on the Internet for the last 16 years. Even proprietary systems (such as Microsoft's Bing search engine) use web interfaces based on open systems. As indicated earlier these systems form the heart of almost all web based e-commerce and social networking systems. They are now being adopted in the most demanding transaction processing environments, such as that provided by Amadeus IT Group the largest transaction processor in the travel industry. Such applications were previously the exclusive reserve of very highly specialised mainframe-based systems. Whilst such a transition from closed proprietary systems to an open approach is both time consuming and challenging (see Chapter Four of this report), there are clear advantages for all stakeholders.

3.2 Advantages to the enterprise customer

3.2.1 Access to greater innovation: The combination of open source software, open systems interoperability and open standards has created a self-reinforcing community of shared research and development and a pooling of creative ideas. This leads to a stream of innovative applications for test and development. Those that show promise are implemented and adopted more quickly.

> **Open source is driving innovation.** In a speech at the Linux leaders annual summit in June 2012, Red Hat CEO Jim Whitehurst commented that the information age is finally evolving into the information economy because of the standardisation enabled by Linux and open source. He said that: "...Open source has gone mainstream ... open source is the default choice of the next generation IT architecture..." and "...More innovation will happen first in open source and that's a radical change from even five years ago..."

3.2.2 Quicker response to changing requirements: The pace of change in many markets, especially those with a consumer focus, is now relentless. The focus on customer engagement and involvement and the influence of burgeoning social networks place further demands on immediacy of response. Access to a wide range of cost-effective development tools, ability to scale rapidly and access to global community of shared knowledge all favour the use of the open approach.

> **Open source's responsiveness.** Alex McLachlan of Indigoblue Consulting shared in his blog that: "One of open source's unique selling points (if I can use that phrase) is its responsiveness. We've had a really good example of this on our website recently, which demonstrates how issues with open source software can get fixed really quickly and efficiently. We use the Drupal open source website content management system and wanted to add in the Amazon module that lets you link directly to Amazon contents. We had an issue with how the links to Amazon were appearing... The guy doing the development posted the problem on the Drupal website and resolved the cause of the problem with the module maintainer within a few days. This experience contrasts with experiences I've had with companies supporting products, where it can be difficult to get fixes made between product releases, so any significant issue can mean waiting until the next product release."

3.2.3 Enhanced ability to support a plethora of systems: The world is a highly heterogeneous environment. Business organisations use a wide variety of current and legacy systems all at different stages of their lifecycle. Some are widely supported, some less so. The ability, through open systems, to draw on a very extensive range of standards and interfaces shared and developed over 20 years can be a key differentiator.

3.3 Advantages for the provider

3.3.1 Access to skilled, motivated and innovative staff: A generation has now grown up with the Internet and open source. They want to work with these systems and tools rather than the closed and proprietary approaches that are frequently seen as more limiting. Open systems developments such as:

- Drupal (an open source content management platform powering millions of websites and applications. It's built, used, and supported by an active and diverse community of people around the world);

- Hadoop (The Apache Hadoop software library is a framework that allows for the distributed processing of large data sets across clusters of computers using a simple programming model. It is designed to scale up from single servers to thousands of machines, each offering local computation and storage. Rather than rely on hardware to deliver high-availability, the library itself is designed to detect and handle failures at the application layer, so delivering a highly-available service on top of a cluster of computers, each of which may be prone to failures); and

- jQuery (a fast and concise JavaScript Library that simplifies HTML document traversing, event handling, animating, and Ajax interactions for rapid web development); are seen as particularly desirable.

3.3.2 Rapid exploitation of new technology developments: Under the continuing impetus of "Moore's Law"[iii] new developments in hardware (such as ultra high resolution displays, enhanced wireless networking and network based storage) continue apace. New developments in software and systems, for example in Business Analytics (the so-called "big data") are also a regular occurrence. It is important to have timely access to these. The open source community once again offers rapid access.

Amadeus IT Group has opened one of its core technology frameworks "Amadeus ARIA™ Templates" to the open source community. This application framework, written in JavaScript, is set to revolutionise the way the travel industry builds new user interfaces, drastically easing the challenges associated with development of travel mobile apps. Third party developers now have the means to propose and build advances to Amadeus' selling interface, opening up a raft of new ideas. As the framework is already used in key Amadeus IT Group products, the company has a long-term commitment to investing in and maintaining ARIA™ Templates.

3.3.3 Drawing on a global community for knowledge tools and problem solving: Developed over the last 20 years, the open source community now numbers in the millions of software architects, analysts, designers and programmers. The community also extends to key universities and facilitates access to leading edge research. There are extensive facilities for collaboration...

GitHub is a social network for open source programmers. It is based around the "Git" the extremely fast, efficient, distributed version control system for the collaborative development of software created by Linus Torvalds. It supports code repositories, issue tracking, collaboration management, code reviews and so on. More than two million software developers are linked through GitHub

openstack™ is a global community of technologists, developers, researchers, corporations and cloud computing experts originally set up under the aegis of the US National aeronautics & Space Administration NASA. It currently involves more than 3000 people from more than 180 organisations. It is based around a massively scalable "cloud" operating system

iii Moore's Law, named after Intel co-founder Gordon E Moore, comes from a 1965statement by Moore that the number of transistors contained in a state of the artintegrated circuit would double every two years. This corresponds roughly to a doublingof processing performance every 18 months.

3.3.4 Limiting restricted dependencies: The open source community avoids dependency on single sources and sharply limits the potential of restricted pricing power. For example, in terms of support for the Linus operating system, competition is based on quality of customer service and additional management facilities rather than on access to the individual system releases. The purchaser has more freedom regarding when to upgrade to a new release...

> **LONDON, February 14, 2011.** Novell today announced SUSE® Linux Enterprise Server will be used to support the London Stock Exchange's new trading platform. The Linux* system will help the Exchange to effectively manage the high volume of traffic passing through its systems in record-breaking time. The deployment, which went live today, will help improve the performance of the Exchange's UK cash markets.
>
> The new Millennium Exchange system will provide the Exchange with a high performance, secure, green and interoperable platform, capable of making intelligent use of virtual and physical environments.
>
> SUSE Linux Enterprise Server will help to increase the Millennium Exchange system's capacity to cope with high-volume trading. SUSE Linux Enterprise Server is a highly reliable, scalable and secure server operating system built to power both physical and virtual mission-critical workloads. With this foundation, enterprises can efficiently deliver business services, enable secure networks and manage heterogeneous IT resources. Millennium Exchange is also being backed by Novell's world-class support and services.

> **May 30, 2012.** Alpha Exchange Improves Performance, Reduces Cost by Migrating to Red Hat Enterprise Linux
>
> Alpha Exchange has been able to keep its position as one of the fastest-trading platforms in Canada through the use of Red Hat® technology. Alpha has seen dramatic improvements in performance, plus cost savings, and enhanced staff productivity.
>
> "The combination of Red Hat Enterprise Linux and HP servers really excels for the type of workloads we do." -Aleksandar Simic, head of infrastructure, Alpha Exchange
>
> Software: Red Hat Enterprise Linux 5.5. Hardware: HP ProLiant Servers
>
> Benefits: Alpha has seen dramatic improvements in performance, plus cost savings, and enhanced staff productivity.

3.3.5 Lower total cost of ownership: Open Source Software is generally free to obtain rather than being subject to the licence fee normally charged for access to proprietary software. However this is not the correct basis on which to compare costs. It is more appropriate to consider the total cost of ownership (TCO) over the lifetime of use of the software. The London School of Economics (LSE) published a report on TCO for the United Kingdom Cabinet Office in 2011. This report found that: "The highest score for strategic drivers was for reduced vendor lock in. A close second was value for money." The report (based on a survey of 32 organisations) found that TCO was often lower overall with open source software. However, based on a model of five life cycle phases (Search for a solution, Acquisition, Integration, Use and Retirement) structured into 14 separate areas of cost, analysis should really be carried out on a case-by-case basis. In so far as generalisation is possible, industry lore is that a TCO saving of around 20% is achievable.

3.3.6 Full visibility of (and confidence in) the source code: There remains no broad agreement in computer security circles as to whether open visibility of source code contributes to, or detracts from, system security. Access to the code is clearly helpful to an attacker, but in the same way extensive public review and correction of the code can minimise the avenues for attack and offer more rapid identification and blocking of successful attacks. These issues are discussed in some depth in the 2002 paper Security in Open versus Closed Systems - The Dance of Boltzmann, Coase

and Moore by Ross Anderson, Professor of Computer Security at Cambridge University. In so far as a consensus exists it is that in pure security terms there is little difference between open and proprietary code. A report by the Communications and Electronic Security Group (CESG) of the UK Government Communications Headquarters (GCHQ) published in December 2011 concluded that: "Open source, as a category, is no more or less secure than closed proprietary software". Trust in "security by obscurity" though does seem misplaced in the proprietary world, as the code is not obscure to insiders. Furthermore, full visibility of code helps to ensure that no "backdoors" or surprises are hidden away...

> "The differences between OSS and proprietary software are not a major factor in either improving or degrading the vulnerability of a nation's IT infrastructure".
>
> Report "Analysis of the impact of Open Source Software" – QinetiQ October 2001

3.4 Advantages to the end consumers

3.4.1 Greater choice and depth of services: Open source, as seen above, both enables greater (as well as faster) innovation and additionally makes it possible to deliver services that would not be financially viable under the economic model of proprietary software.

3.4.2 Democratisation of services: Functionality that would otherwise only be viable to provide on a limited scale or at a high price point – as it would be supported either by proprietary software or by manual activity – can be offered widely when supported by open source software. By providing an operating model without volume-related licence fees, open solutions are more feasibly scalable – hence services that previously could only be offered to top end consumers can be made available to a broader population. This in turn allows yet greater choice, more informed decision making and a better consumer experience.

Simply put, many of the services based on open source – from search engines to social media to holiday planning – would either never have been developed, or could not feasibly be operated, had they not been developed on open source.

3.5 Governance and audit

3.5.1 The challenge of tightly coupled systems: The last 20 years have seen the global economy become ever more interconnected and efficient. Organisations have been encouraged to run lean, become flatter; to take out stock and to operate 'just in time'. Outsourcing, and focusing only on core competencies, has become a mantra, trading off cost reduction against reduced visibility and control. Much capital previously tied up in stock or 'work in progress' has been liberated for more effective investment. In terms of competitiveness, all of this makes absolute sense, and is, indeed, essential for economic success. However, it has become clear that there is a ghost hovering at this particular feast. Running lean makes excellent business sense, but it contains an implicit assumption: that the key systems must be secure and reliable. The World Economic Forum gives pride of place to risks from IT systems on its map of the global risk landscape – see Figure Two.

Figure Two: the global risk landscape: technological risks

With grateful acknowledgement to the World Economic Forum.

3.5.2 The new challenge of audit in the cloud: In order to assess the integrity of IT systems, auditors traditionally rely upon two vital characteristics. These are control and visibility. In cloud operation, elements of these move to the cloud provider, dependent on the model of cloud adopted (see chart in section 2.4.2). These are elaborated in the report Cloud Computing Synopsis and Recommendations published in February 2012 by the US National Institute of Standards.

- "Control: the ability to decide, with high confidence, who and what is allowed to access consumer data and programs, and the ability to perform actions (such as erasing data or disconnecting a network) with high confidence both that the actions have been taken and that no additional actions were taken that would subvert the consumer's intent (e.g., a consumer request to erase a data object should not be subverted by the silent generation of a copy).

- Visibility: the ability to monitor, with high confidence, the status of a consumer's data and programs and how consumer data and programs are being accessed by others".

IT development in general and cloud computing in particular should be high on the agenda of corporate audit committees. The KPMG 2012 Audit Institute Report identified "IT Risk and Emerging Technologies" as the second-highest concern for audit committees, behind "Governance Processes, Controls and Risk Management."

"IT Risk" had been ranked sixth in 2011.

3.5.3 Old models may be obsolescent: In the world of the wholly owned corporate data centre, prior to virtualisation, control was clear and retained fully within the organisation itself. Visibility was straightforward, particular physical machines ran well-defined applications, data was stored in well understood multiply redundant databases again in well defined locations on specific hardware. It was possible literally to audit physically. Much of this changes with the advent of the cloud. Many responsibilities now move either wholly or partially to the cloud provider including areas such as security policy, asset management, human resource security, physical and environmental security, access controls, configuration management, business continuity management, security incident management, legal compliance, etc. Particular applications are no longer tied to specific hardware or even specific locations. This may become especially problematic if (say under failure recovery conditions) data sets are moved across jurisdictional boundaries (for example out of European Economic Area to locations not covered by "EU-US agreed Safe Harbor" principles).

The European Data Protection Directive: There is a well-developed legal framework in the European Union (EU) with respect to the protection of personal privacy. Transfer of personal data to countries outside the EU/EEA is subject to regulation. The main principle is that such data should only be transferred to jurisdictions that provide suitable levels of comparable protection. Many jurisdictions are regarded as not offering adequate protection. It should be noted that new draft European Data Protection Regulation is currently being prepared. There are also concerns about access by Governments to such data under national legislation such as the US Patriot Act.

3.5.4 New approaches are required: In the world of the cloud, audit will need to rely far more on assessment of the stated policies and controls (for example in accordance with the ISO 27000 series standards), plus spot checks carried out by the cloud providers' own auditors. The "SSAE 16"[iv] report from the cloud provider's auditors will be particularly important. Public cloud providers in general will not accept requests for third party audit. There will need to be extensive cooperation between the auditors employed by those contracting services out to the cloud and the auditors for the cloud providers. The contractual service level agreements are complex and must focus on how to cope if things go wrong for example how (and how quickly) will applications and data be retrieved if the cloud provider were to go into liquidation... These developments are not optional. They are fundamental to ensuring the continued integrity of key systems.

3.5.5 The role of professional bodies: It will be important for the accounting profession, and oversight bodies such as the UK Financial Reporting Council, to demonstrate how the necessary audit and risk management assessments can now be appropriately carried out and trust fostered? This will be challenging when key business processes are delivered through cloud solutions that are not necessarily locked to specific hardware or to geographical locations.

Quick take:

Major and demonstrated benefits accrue to all stakeholders from a migration to open systems:

- Key advantages for enterprise customers of the move to open systems typically include: access to greater innovation; improved supplier responsiveness; and enhanced system accessibility.
- Key advantages for providers include: lower total cost of ownership; drawing on a broader pool of highly skilled and motivated staff; quicker access to new technology developments; access to a global community for problem solving; reducing monopoly dependencies and full confidence in the visible source code.

However, the final move to cloud computing will demand new approaches in risk management and audit and close collaboration between the auditors of service users and those of the service providers. There are significant challenges along the way as highlighted in the next chapter...

iv Statement on Standards for Attestation Engagements No. 16 has recently replaced the Statement on Auditing Standards No. 70 (SAS 70). These are widely recognised auditing standards developed by the American Institute of Certified Public Accountants (AICPA). A service auditor's examination performed in accordance with SSAE 16 or SAS 70 is widely recognised, because it represents that a service organisation has been through an in-depth audit of their control objectives and control activities, which often include controls over information technology and related processes.

4. Lessons from the transition from closed to open systems

4.1 As the path away from dependency on proprietary systems becomes evermore well worn, it is helpful to document the "lessons learnt" by those who have already made, or are deeply engaged, in that transition. There is no value in re-visiting pitfalls already explored.

4.2 Maintaining a common development/support team across existing and replacement systems: the key reasons for maintaining common teams are:

- allowing the staff that have been responsible for supporting the old systems to be trained in the new programming languages, methodologies and toolsets, thus demonstrating commitment to avoiding them becoming a "stranded asset" as the old systems are progressively retired.

- maintaining the "institutional memory" of why particular design and implementation choices were made, these may still be relevant to the new developments.

- giving development teams a role in support as well, thus ensuring that they are fully aware of, and live with, the consequences of implementation defects.

4.3 Accepting the challenge of the task and being realistic about the costs, manpower and timescale required: this is a major undertaking. Often this represents a significant business change, not simply an IT change – see figure three overleaf. New systems must be developed whilst maintaining high levels of service to customers through the existing systems. In the words of Henry Kissinger: "...the urgent often forces out the important". Everyone was busy before, now there are lots of extra change related tasks as well. Emotions rise as workloads increase and people are inexorably sucked into the imperatives of day-to-day business rather than fully addressing the change.

4.4 Recognising that it is more than replicating existing functions: a careful analysis is required to:

- weed out requirements that are simply no longer relevant.

- carry out needed rationalisation and simplification of requirements changes that have accumulated in the old system over years; and

- include new features and capabilities to support anticipated customer.

- Despite the best of intentions to minimise change, it may simply not be possible to freeze all new development requirements for the duration of the change programme. Some degree of revision during the development process may be inevitable.

4.5 Significant changes in operational procedures will need to be planned and tested including:

- Ensuring data integrity during rollback/recovery: Where elements required as part of a specific transaction are distributed across multiple databases, whether on specified servers or within a cloud, special measures are required to ensure data integrity is maintained if the transaction is aborted or some element is subject to failure. This can represent a significant challenge for example in rolling back such a system to a known prior state. The traditional database properties of "atomicity", "consistency", "isolation" and "durability" (ACID) all need careful attention in the context of maintaining integrity in the presence of aborted or failed transactions.

- "Seeing the wood for the trees" in fault management: In distributed or cloud based systems a single initial failure can lead to a rapidly branching tree of subsequent errors and alarms. Expert system techniques may be required to isolate rapidly the root cause otherwise hidden in the noise of cascading failures.

Figure Three: what is the most appropriate budget split between "technology" spend and "people" spend in major business change?

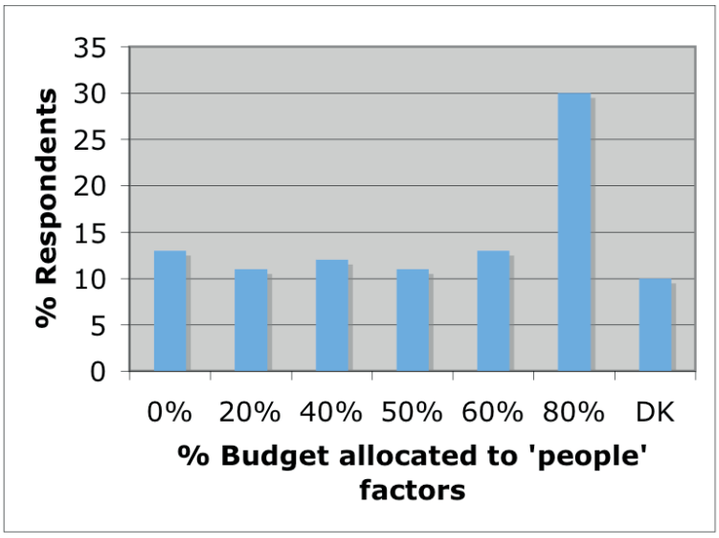

'People' spend includes: business process re-engineering, internal communications, training, changes to performance measurement, reward....

Source: IoD Business Opinion Survey researched May 2006

Quick take:

Whilst very much worthwhile, the path to open systems is challenging. It is important to learn from the experience of others and so to navigate around the known pitfalls:

- Don't separate out support for the current systems from development of the new.
- Be realistic at the outset about costs, manpower and timescale – it is neither quick nor easy.
- Only transfer what you really need and absolutely minimise new development during the transfer.
- Ensure that your operational procedures reflect the extra work required in both data integrity during rollback/recovery and in managing fault trees.

5. What might the future of OSS and cloud computing hold?

5.1 The physicist Niels Bohr pointed out that: "Prediction is very difficult, especially about the future". Undeterred, it is possible to do some forward projection based on work already in progress today.

5.2 Ubiquity or efficiency? The current structures of cloud computing provide immense utility. They have four key technical characteristics, well described in the paper Cloud Computing Technologies from the Intelligent Systems Research Centre at the University of Ulster:

- multi-tenancy – resource pooling allows many customers to draw dynamically on resources from a central pool.
- massive scalability – users can self-provision resources such as server or storage capability, within the bounds of a contract, as demand requires.
- rapid elasticity – users can very quickly increase or decrease the allocated resources.
- measured service – the user only pays for the resources actually consumed.

Such distributed parallel computing has been made possible by models such as MapReduce, which hide this complexity and the need to manage synchronisation, data motion and communication, behind a simple interface. They underpin the OSS such as Hadoop[v]. Such ubiquity can mask very substantial inefficiency. This is described comprehensively in the paper "The seven deadly sins of cloud computing research". It seems likely that new and more efficient approaches will be required before very high performance transaction processing systems can draw on either private or public cloud resources. Even for a conventional OSS software stack such as LAMP[vi] there is substantial cost of complexity and inefficiency. This is summed up in a cartoon from the web comic xkcd:

A new approach will be required for very high performance cloud processing. Such work is already in progress. One example is research on the cloud operating system "Project Mirage". This aims to offer simplicity, efficiency and security through compiling applications directly into a high-performance, energy-efficient kernel. The conventional seven layers of hardware, Hypervisor, OS Kernel, User Processes, Language Runtime, Threads and Application Code are condensed into just four Hardware, Hypervisor, Application Code and Mirage Kernel. For a further description please see the paper Multiscale not Multicore: Efficient Heterogeneous Cloud Computing. This approach offers the promise of scaling with equal ease across multiple cores in a single processor or multiple processors in a public or private cloud.

v The Apache Hadoop software library is a framework that allows for the distributed processing of large data sets across clusters of computers using a simple programming model. It is designed to scale up from single servers to thousands of machines, each offering local computation and storage. Rather than rely on hardware to deliver high-availability, the library itself is designed to detect and handle failures at the application layer, so delivering a highly available service on top of a cluster of computers, each of which may be prone to failures.

vi Software Stack is a term of art for example a typical cloud computing stack would be Linux (operating System), Apache (web server), MySQL (open source database) and PHP (scripting language) – known colloquially as "LAMP"

5.3 Hybrid cloud: a desirable option for many organisations would be to operate their applications through a private cloud fully under their own control, but with the option to draw on the services of a public cloud under overload or partial failure conditions. Tools to facilitate this are now becoming available. One example is the "vCloud Connector" from VMware.

5.4 Federation between clouds: a key interest for those using either public or private clouds will be interoperability between different cloud providers and portability of data and applications between cloud providers. At present this is very challenging. Common interfaces (APIs) have yet to be properly defined. The European Commission report: Advances in Clouds: Research in Future Cloud Computing has some trenchant observations in this area.

5.5 Neutrality: the future looks likely to entail an inexorable shift to applications that are "operating systems neutral" and thus not locked to any specific operating system or browser. Figure 4 overleaf shows a projection by Joe Baguley, Chief Technologist at VMware. It is contained in the report The future of computing: indispensable or unsustainable?

Figure Four: graph showing how the number of Windows-specific applications being developed has decreased in favour of OS neutral apps.

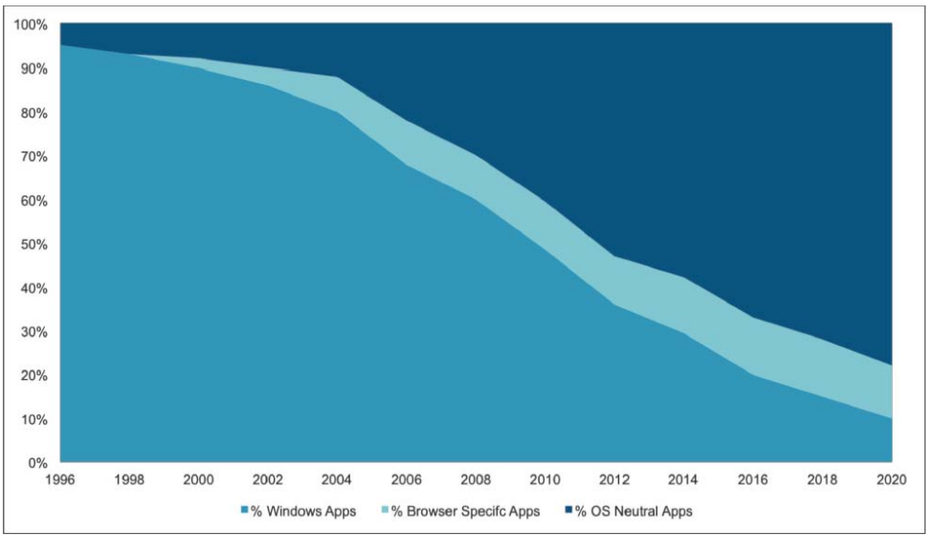

Quick take:

The story of OSS is very much one of continuing and rapid development. Research is in progress to:

- Minimise the efficiency penalty that comes with the ubiquity of current generation systems.
- Ensure that it is both possible and economic to operate a hybrid (partly private, partly public) cloud solution.
- Develop interoperability between public cloud providers ensuring portability of both applications and data.

6. Conclusions

This white paper has sought to place the 40 years of development of open source Software and open systems into its historical context. These developments have now reached the point of maturity where they are appropriate for use in even highly demanding business critical applications. The real benefits have been demonstrated. The next stage, a transition to virtualisation and thence onwards to both private and public cloud computing, looms for many organisations. Here there are still challenges to be overcome for high volume, business critical, applications. It will be important to minimise the inefficiencies inherent in the current open source software stacks such as LAMP[vii] before this transition can be made and a final tranche of benefits for all stakeholders released. However, promising lines of research exist and it is clear that these challenges can, and will, be overcome.

vii Linux (operating System), Apache (web server), MySQL (open source database) and PHP (scripting language) – known colloquially as "LAMP"

7. References

i GNU is an open source operating system

ii Transmission Control Protocol/Internet Protocol

iii Moore's Law, named after Intel co-founder Gordon E Moore, comes from a 1965 statement by Moore that the number of transistors contained in a state of the art integrated circuit would double every two years. This corresponds roughly to a doubling of processing performance every 18 months.

iv Statement on Standards for Attestation Engagements No. 16 has recently replaced the Statement on Auditing Standards No. 70 (SAS 70). These are widely recognised auditing standards developed by the American Institute of Certified Public Accountants (AICPA). A service auditor's examination performed in accordance with SSAE 16 or SAS 70 is widely recognised, because it represents that a service organisation has been through an in-depth audit of their control objectives and control activities, which often include controls over information technology and related processes.

v The Apache Hadoop software library is a framework that allows for the distributed processing of large data sets across clusters of computers using a simple programming model. It is designed to scale up from single servers to thousands of machines, each offering local computation and storage. Rather than rely on hardware to deliver high-availability, the library itself is designed to detect and handle failures at the application layer, so delivering a highly available service on top of a cluster of computers, each of which may be prone to failures.

vi Software Stack is a term of art for example a typical cloud computing stack would be Linux (operating System), Apache (web server), MySQL (open source database) and PHP (scripting language) – known colloquially as "LAMP"

vii Linux (operating System), Apache (web server), MySQL (open source database) and PHP (scripting language) – known colloquially as "LAMP"

1 See Open Source Initiative at http://opensource.org/osr

2 Source: UK Houses of Parliament Postnote 414. See: http://www.parliament.uk/business/publications/research/briefing-papers/POST-PN-414

3 OSI Reference Model-The ISO Model of Architecture for Open Systems Interconnection, Hubert Zimmermann, IEEE Transactions on Communications, Vol. Com-28, No. 4, April 1980.

4 For a history of Unix, see: http://www.levenez.com/unix/

5 For a description of Linux see: http://en.wikipedia.org/wiki/Linux

6 For the full text from the Google blog please see: http://googleblog.blogspot.co.uk/2009/12/meaning-of-open.html

7 To download the IET Fact file go to: http://www.theiet.org/factfiles/it/cloud-computing-page.cfm

8 Gartner research quoted in Computer Weekly. See: http://www.computerweekly.com/news/2240161014/Global-IT-outsourcing-spend-to-exceed-250bn-in-2012

9 Source Apache Software Foundation. See: http://httpd.apache.org/

10 See: http://www.zdnet.com/blog/open-source/red-hat-open-source-is-driving-innovation-and-the-information-economy-but-battle-is-not-over/11342

11 See: http://indigoblue.co.uk/strategic-consulting/blog/open-sources-responsiveness

12 See: http://drupal.org/

13 See: http://hadoop.apache.org/

14 See: http://jquery.com/

15 See: http://www.amadeus.com/amadeus/x224333.html

16 See: https://github.com/features/community

17 See: http://www.openstack.org/

18 http://www.suse.com/company/press/2011/2/novell-delivers-linux-server-system-to-london-stock-exchange.

19 Shaikh, Maha and Cornford, Tony (2011) Total cost of ownership of open source software: a report for the UK Cabinet Office supported by OpenForum Europe. UK Cabinet Office, London, UK. See: http://eprints.lse.ac.uk/39826/

20 Anderson, R. J. (2002). Security in open versus closed systems---the Dance of Boltzmann, Coase and Moore. Presented at Open Source Software Economics. See: www.cl.cam.ac.uk/~rja14/Papers/toulouse.pdf

21 UK Cabinet Office – Open Source Software Security December 2011.

22 Peeling, Nic and Satchell, Julian (2001) Analysis of the Impact of Open Source Software. See: www.cabinetoffice.gov.uk/media/253425/QinetiQ_OSS_rep.pdf

23 Source: World Economic Forum Global Risks 2012. See http://www.weforum.org/reports/global-risks-2012-seventh-edition

24 Badger, M. L.; Grance, T.; Patt-Corner, R.; Voas, J. M. (May 2012) Cloud Computing Synopsis and Recommendations. See: http://www.nist.gov/manuscript-publication-search.cfm?pub_id=911075

25 Is Governance keeping pace? Challenges and Priorities shaping the Audit Committee Agenda. Available from: www.kpmg.com/Global/en/IssuesAndInsights/ArticlesPublications/Documents/kpmg-2012-audit-committee-issues-conference.pdf

26 See: http://www.justice.gov/archive/ll/highlights.htm

27 Carlin, Sean and Curran, Kevin (June 2012) Cloud Computing Technologies, Journal of Cloud Computing and Services science Vol. 1, No. 2. http://iaesjournal.com/online/index.php/IJ-CLOSER/article/view/Northern%20Ireland%2C%20UK/pdf

28 DEAN, J., ET AL. MapReduce: Simplified Data Processing on Large Clusters. In Proc. OSDI (2004), p. 10.

29 Schwarzkopf, Malte; Murray, Derek; Hand, Steven. The seven deadly sins of cloud computing research. Proceedings of the 4th Workshop on Hot Topics in Cloud Computing (HotCloud 2012) at USENIX 2012. See: http://www.cl.cam.ac.uk/~ms705/pub/

30 See: http://www.xkcd.com/676

31 See: http://www.openmirage.org

32 Madhavapeddy, Anil; Mortier, Richard; Crowcroft, Jon; Hand, Steven. Multiscale not Multicore: Efficient Heterogenous Cloud Computing. Proceedings of Visions 2010. See: www.cl.cam.ac.uk/~jac22/out/orizens.pdf

33 See: http://www.vmware.com/products/datacenter-virtualization/vcloudconnector/overview.html

34 The report can be downloaded from: cordis.europa.eu/fp7/ict/ssai/docs/future-cc-2may-finalreport-experts.pdf

35 See: www.cl.cam.ac.uk/~acr31/pubs/raeng-futurecomputing.pdf

amadeus

Your technology partner