

Grid Computing Update The Market Gathers Momentum

Contacts:

Bob Tarzey
Quocirca Ltd
Tel +44 1753 855794
bob.tarzey@quocirca.com

Clive Longbottom
Quocirca Ltd
Tel +44 118 948 3360
clive.longbottom@quocirca.com

In the third quarter of 2004, Quocirca reported on the state of Grid Computing adoption in Europe through a set of key market indicators that collectively made up the "Grid Index". In March 2005, these same indicators provide clear insights into the way Grid related activity is developing, but this time on a global basis.

- **IT professionals are now much better informed on the nature and benefits of virtualisation and grid related technologies**

Trending data for Europe, where two cycles of the Grid Index programme have now been executed, indicate a doubling in the level of knowledge of virtualisation and grid related technologies over the past 9 months. The key indicator in this area, the "Knowledge Index", increased from 2.74 in June 2004 to 5.81 in March 2005 (on a scale of 0 to 10). There was a similar dramatic increase in the "Benefit Index", an indicator of how much IT professionals acknowledge the potential of the technologies tracked. This index rose from 2.25 to 4.58 over the same period.

- **An increase in marketing, sales and education by IT vendors has driven the awareness**

The main source of Grid Computing knowledge cited by survey respondents was the IT vendor community, with Web communities and special interest groups also contributing significantly. The media have yet to ramp up their activity in this area, though we are starting to see an increase in coverage by both the business and IT press.

- **Virtualisation technologies are gaining significant mainstream traction, which in turn will drive Grid Computing adoption**

The "Commitment Index", which measures the degree to which organisations are incorporating emerging solutions into their plans and activities, rose by between 50% and 80% for server virtualisation and storage virtualisation technologies between June 2004 and March 2005. This suggests that these technologies are now moving freely into mainstream adoption. Taking virtualisation to the next step by introducing automation in the form of Grid Computing is still only taking place amongst a relatively small number of early adopters, however, with modest growth (<15%) in tangible activity off a smaller base. Grid Computing therefore still has to move from idea to action for most organisations, but such a move is very likely as the momentum behind virtualisation will drive this natural evolution.

- **Adoption in North America is ahead of Europe, and the Asia Pacific region lags behind**

The combined Commitment Indices for virtualisation and Grid technologies in North America, Europe and Asia Pacific were 3.02, 2.70 and 1.74 respectively. The highest local adoption rates worldwide were recorded in the Nordic region of Europe, which had a Commitment Index of 3.60. The lowest adoption rates were in China, Korea and India. This distribution of early activity is normal for emerging technologies. However, there is clear evidence that a higher level of infrastructure standardisation in Asian countries increases the appeal of Grid Computing and the ease with which it can be implemented. Asia Pacific is therefore well positioned to close the gap.

- **Grid adoption goes hand in hand with modern software and hardware architectures**

There is a clear correlation between the knowledge and adoption of grid related technologies and familiarity with and use of service oriented architectures (SOA). This makes sense as the componentised application approach associated with SOA enables more flexibility and control in a grid environment, enhancing the benefits significantly. Those more active with grid related technologies were also significantly more likely to be using blade computing architectures.

- **Whilst challenges remain, the positive steps towards grid are very encouraging**

Despite the dramatic increases in awareness, IT professionals still highlight limited familiarity and skills availability as two of the main constraining factors on grid computing adoption, along with the need for solutions to mature further and for standards to firm up. Given that the industry has moved up a gear in addressing these issues and early experience with virtualisation will drive interest in automation, the outlook for the mainstream grid computing market is very healthy.

RESEARCH NOTE:

The information presented in this report was derived from 1356 interviews with senior IT influencers and decision makers completed in Mar 2005. 300 of these were conducted in North America, 600 in Europe and the remainder (456) in Asia Pacific. Respondents were from a mixture of large multinationals and medium to large national organisations.

CONTENTS

| | | |
|----|---|----|
| 1. | INTRODUCTION..... | 3 |
| 2. | GRID AS AN EVOLUTIONARY CONCEPT | 3 |
| 3. | PRACTICAL CONSIDERATIONS..... | 3 |
| 4. | CURRENT LEVEL OF PROGRESS AND ACTIVITY | 4 |
| 5. | MARKET DYNAMICS..... | 4 |
| 6. | RELATIONSHIP BETWEEN GRID AND OTHER TECHNOLOGIES..... | 6 |
| 7. | CONFIRMATION OF GRID BENEFITS..... | 6 |
| 8. | THE APPLICATION OF GRID..... | 7 |
| 9. | DISCUSSION | 7 |
| | APPENDIX A - A BUSINESS VIEW OF IT INFRASTRUCTURE EVOLUTION | 8 |
| | STANDARDISATION..... | 8 |
| | SERVER CONSOLIDATION | 8 |
| | POOLING AND VIRTUALISATION..... | 8 |
| | GRID COMPUTING | 9 |
| | APPENDIX B – ORIGIN OF THE GRID INDICES..... | 11 |
| | APPROACH AND RATIONALE..... | 11 |
| | DATA GATHERING AND INDEX CALCULATION | 11 |
| | APPENDIX C – INTERVIEW SAMPLE DISTRIBUTION..... | 12 |
| | ABOUT ORACLE | 13 |
| | ABOUT QUOCIRCA | 14 |

1. Introduction

The IT industry is full of visionaries. It sometimes seems like every way we turn we run into visions from suppliers showing us how things will be in 5, 10 or 20 years' time.

One of these visions paints the picture of a world in which computing power is delivered as a utility service to anyone that wants it, in the same way as electricity is delivered today via a socket in the wall. This is an interesting concept, but when we look at the highly specified PC in our home or the rack of servers in the computer room at work, we have to wonder how long it will be before the vision becomes a reality. The chances are that it will be some time before all of that local computing power is replaced by a cable running into our homes or businesses from a central source, via a network akin to an electricity supply grid.

And this brings us to one of the biggest challenges with so called grid computing – the misperception of what it's for. Articles in business journals often position it as an enabler of the futuristic "computing-on-tap", which makes it sound far removed from the day-to-day activities of the world we currently live in. Matters are made worse when we read stories of how grid computing is used in the search for extraterrestrials, aligning it with science fiction.

As a result there has been a tendency for grid computing to be thought of as a possible technology to watch for the future, rather than as something of practical relevance today.

The Grid Index programme upon which this report is based was designed to look at grid computing from another angle. Rather than working backwards from a vision, we started with the real world of mainstream business computing as it is today and considered how grid computing might build on that. When dealt with in this way, it is clear that grid is simply a step in the natural evolution of IT infrastructures.

2. Grid as an Evolutionary Concept

The view of grid computing as an evolutionary concept was documented fully in the last Quocirca Grid Index report published in September 2004 (www.quocirca.com/report_grid_2004.htm). For convenience, the background chapter from this report entitled "A Business View of Infrastructure Evolution" is reproduced in Appendix A. Those unfamiliar with the concepts of virtualisation and grid computing may use this as a primer or reference.

In summary, though, we can consider that mainstream organisations typically go through a number of logical steps as they look to optimise different parts of their IT infrastructure. The first step, if it has not been done already, is to remove excessive diversity by standardising applications and platform components as much as possible. As this is done, it is common to create or reveal redundancy between components which then leads quite naturally to a degree of consolidation.

This type of standardisation and consolidation activity is usually justified because it drives efficiency. There is less need for diverse specialist skill sets and IT operations are generally simplified, both of which translate to a lowering of cost and business risk.

But introducing more consistency into the infrastructure also leads to another important benefit. If many of the elements in the network now look the same or similar, it is much easier to use components together in a coordinated manner. It is also

easier to add or remove components and re-allocate them to different tasks as requirements change.

This is the basis for yet more optimisation potential based on the introduction of virtualisation technologies and grid computing. The use of such technologies leads to more flexibility and responsiveness to business needs as the IT department is able to react more quickly to changing requirements and demands.

There are also further cost savings, achieved through boosting hardware utilisation levels, i.e. by unlocking more of the inherent power of computer systems to do useful work.

Introduction of "grid technologies" is thus the next logical step for those that have achieved a degree of consistency in their IT infrastructures through standardisation and consolidation.

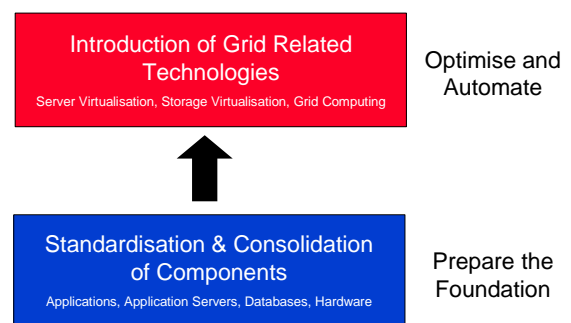
The Grid Index programme was designed to monitor the rate at which this kind of evolution is taking place across the enterprise sector.

3. Practical Considerations

In order to monitor the way infrastructures are evolving, it is useful to consider two activities as illustrated in Figure 1.

Figure 1

Driving Efficiency, Flexibility and Responsiveness



Of course in the real world, these are not two distinct waves that roll across an organisation's entire IT infrastructure in sequence. In practice, these steps take place as a series of parallel initiatives in different parts of the business, often at different times. Sometimes activity is coordinated as part of an overall strategy, but sometimes it takes the form of a number of completely separate tactical exercises. Discrete optimisation initiatives in relation to ERP systems, Web transaction systems, email systems and so on, are quite common.

We should also bear in mind that complete standardisation is not possible in most larger-scale environments. It would be nice, for example, to move all applications onto either a J2EE or Microsoft .Net application server for the sake of operational simplicity and efficiency. However, this is a luxury available to few, as legacy and packaged application dependencies typically dictate that multiple operating systems, databases, application servers and other platform components are maintained.

When we talk about standardisation in the context of overall infrastructure evolution, we are therefore referring to a compromise between the relative chaos of highly fragmented and disjointed systems that arise from piecemeal organic growth, and the ideal situation of a completely homogenous


environment. The reality is that optimisation and automation solutions are generally easier to implement and more efficient to run if there is a higher degree of infrastructure consistency.

When designing the Grid Index programme as a monitoring tool, it was necessary to take into account all of these real-world considerations, particularly the often parallel and iterative nature of activities as we have described. Foundation preparation is therefore measured separately to the adoption lifecycle associated with the introduction of optimisation and automation solutions.

An explanation of how the Grid Index is designed is provided in Appendix B. For the purposes of interpretation, all we need to know is that what we end up with is a set of indicators that tell us how strong activity or progress is in each of the relevant areas.

4. Current Level of Progress and Activity

The following picture provides a summary level snapshot of activity and progress on a global basis (Figure 2):



The results upon which this picture is based were derived from 1356 interviews with senior IT influencers and decision makers from around the world completed in March 2005 (see Appendix C for sample distribution).

We can infer from the picture that foundation readiness is at a reasonable level - certainly high enough in the majority of organisations to permit the introduction of grid technologies into at least some parts of the infrastructure.

We can also see that IT professionals have a good degree of knowledge and understanding of how grid technologies can be of benefit in business terms. The Benefit Index measures the perception of value associated with these solutions enabling operational efficiency improvements and allowing a better and more responsive service to the business.

Commitment levels are based on the degree to which grid technologies have been incorporated into the IT strategy, current levels of adoption and planned future investment. As we can see, the indicator that measures this, the Commitment Index, is significantly lower. This tells us that knowledge and interest has yet to translate into tangible plans and activity in most cases. The ROI Index, which is based on achieved or expected return on investment, is not surprisingly also quite modest.

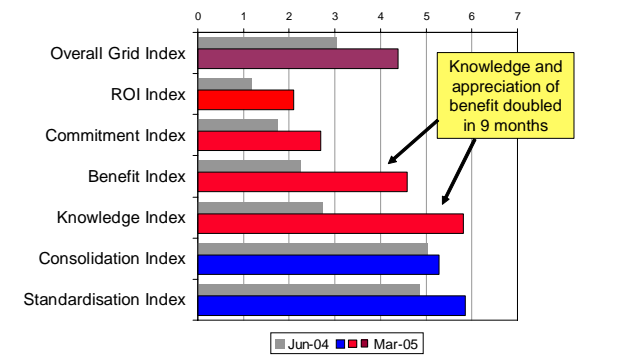
But this snapshot view only tells part of the story.

Whilst March 2005 was the first time this research was conducted in North America and the Asia Pacific region, the Grid Index programme started life in the middle of 2004 as a European initiative. This means we have trending data for Europe that provides us with some insight into how quickly interest and activity in grid technologies is building.

5. Market Dynamics

In the nine months between the first and second cycle of Grid Index research in Europe, some dramatic increases were observed, especially in the areas of awareness, knowledge and appreciation of the benefit of grid technologies (Figure 3).

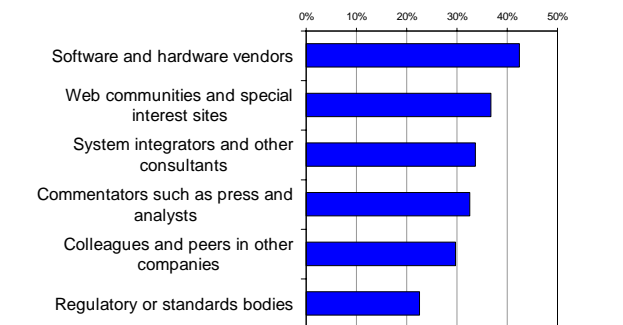
Figure 3
European Trends Indicate Rapid Building of Momentum



European IT professionals, are not only becoming much more informed, but are also starting to appreciate how these technologies will drive more value out of their IT systems (it should be noted that the individuals that took part in the 2005 survey were not the same as those from 2004, so the increase in knowledge is not driven by the research itself).

The most important driver of this increased knowledge and interest is the education, marketing and sales activities of IT vendors (Figure 4).

Figure 4
Where has your knowledge and awareness of Virtualisation and Grid Computing technologies primarily come from?



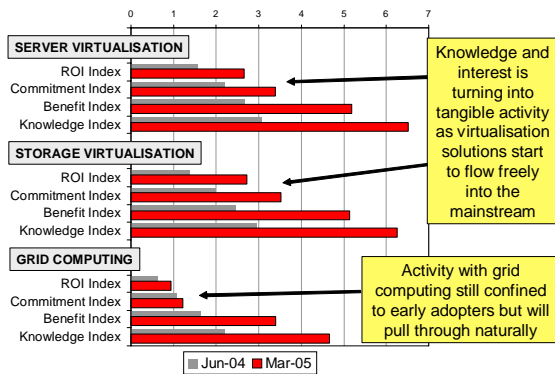
This emphasis on suppliers is quite natural at this stage as vendors keen to “make the market” are beginning to put much more effort into messaging and education, the levels of which have been ramped up considerably over the past 6 to 9 months.

The prominence of Web communities and special interest groups is also very encouraging. One of the constraints on moving forward with new technologies is often a lack of experience. Peer based communities are an important part of working around this barrier.

As yet, the press and analyst community have generally not moved much beyond discussion of grid computing as a curiosity, as we discussed in the Introduction. Unlike some of the more hyped areas of the industry, it therefore appears that many IT professionals in the mainstream are actually ahead of the pundits and commentators in appreciating the relevance of grid. We anticipate that this will change as more journalists and analysts begin to realise the importance and relevance of grid to their readers and subscribers.

Whilst we are still in what many would call the “early market”, all of the signs indicate that the mainstream is set to embrace and reap the benefits of grid related technologies. Indeed, some of the components of grid computing, namely the virtualisation technologies, are already gaining traction in terms of adoption (Figure 5).

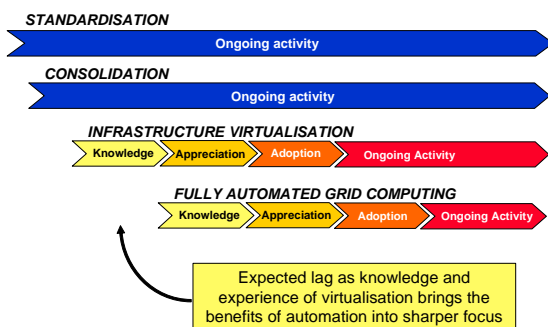
Figure 5
Virtualisation Technologies Gaining Traction Rapidly



Whilst this is a strong indication of mainstream acceptance of server and storage virtualisation, this picture seems to suggest that there is a problem with moving from theory into practice with grid computing, which adds full automation to virtualised infrastructures (please read Appendix A if some of these concepts are unfamiliar).

However, a lag is to be expected between organisations taking the infrastructure virtualisation concept on board and moving on to the next stage of automation (Figure 6).

Figure 6
Evolution of Ideas and Activity in Practice




There is also the practicality of organisations not trying to run before they can walk. Attempting to implement fully automated grid computing without first gaining experience of manually configured and managed virtualisation solutions will not make sense in the majority of cases. Even IT vendors with capable automation technology recommend against this. The only exception to this might be turn-key implementations by professional services firms who subsequently continue to manage the environment or go

through a controlled hand-back with the relevant education and hand-holding worked into the process.

Another interesting aspect of market dynamics is the variation of activity by geography. When we examine this dimension, we find, not surprisingly, that North America is ahead of Europe in terms of tangible commitment with the Asia Pacific region lagging (Figure 7).

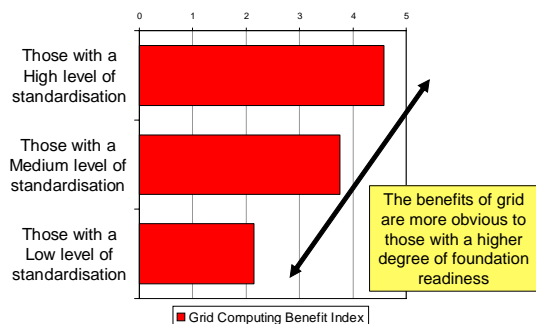
Figure 7
Regional Variation



As noted on the figure, however, a higher degree of foundation readiness was observed in the Asia Pacific region, particularly in areas like South East Asia. This makes sense when we consider the recent waves of technology modernisation that have swept these countries along with local language requirements acting as a natural constraint on diversity. The combination of these factors means that the average Asian IT infrastructure is inherently a lot less diverse and complex than the average American or European one.

This is significant when we consider the relationship between standardisation and appreciation of the benefits of grid technologies (Figure 8).

Figure 8
Relationship between the Level of Standardisation and the Appreciation of Benefits



Based on this, it could be argued that the higher foundation readiness observed in the Asia Pacific region could accelerate the adoption of grid technologies and close the gap with North America and Europe. This potential is reflected in the fact that the overall grid index, which takes into account platform readiness as well as the grid technology adoption lifecycle, does not vary that much by region:

| | |
|---------------|------|
| North America | 4.50 |
| Europe | 4.39 |
| Asia Pacific | 4.37 |

6. Relationship between Grid and Other Technologies

The correlation between the level of standardisation and the appreciation of benefits illustrated in Figure 8 is an example of how analysis of the Grid Index research may be used to highlight relationships and dependencies.

We used this same technique to look at the relationship between commitment to grid and the use of other technologies. This revealed that those with a higher commitment to grid were also more inclined to be using modern hardware and software architectures in the form of blade servers (Figure 9) and service oriented architectures (SOA, Figure 10).

Figure 9

Do you make significant use of Blade Server Architectures today or are you likely to do so in the next 12 months?


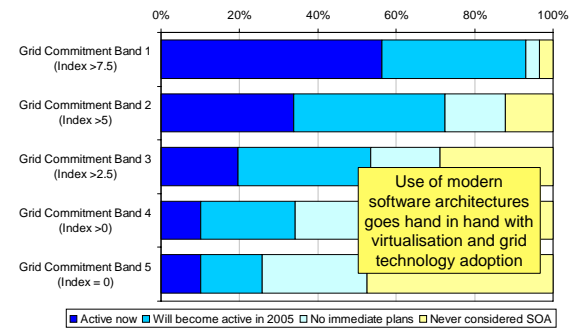


Figure 10

Those active or committed to activity with Grid are more likely to be active with SOA




Both of these correlations make sense as Blades, SOA and Grid are all aimed at flexible delivery of IT. In terms of the physical infrastructure, Blade Computing allows hardware components to be added and removed efficiently with minimal disruption, and for processing power, memory etc to be maximised depending on requirement. The component model that underpins SOA means software applications are delivered as smaller operational units that may be started up, shut down and distributed in a much more precise, flexible and controllable manner compared to traditional monolithic applications. And, of course, grid technologies deal with the relationship between hardware and software, providing flexibility by managing which software components are serviced by which hardware components on a dynamic basis.

Whilst there is no absolute dependency between these three technologies, i.e. it is perfectly possible to implement any of them separately to the others and still achieve significant benefits, incorporating all of them together into an IT

strategy is a way of maximising the overall level of flexibility of a given systems environment (Figure 11).

Figure 11

Exploitation of Synergies



Following on from this, the more knowledgeable respondents highlighted the need for packaged application vendors to move their solutions onto an SOA footing if they haven't done so already, in preparation for use in a grid environment.

Regardless of the use of other technologies, however, there is clear confirmation from the research of the specific benefits of grid solutions, beyond just flexibility, into cost savings and responsiveness to the business.


7. Confirmation of Grid Benefits

As we have already discussed, grid computing theory tells us that we should expect better hardware utilisation across an IT infrastructure. The ability to move available resource around in response to peaks and troughs in application demand means there is less need to "oversize" individual machines to deal with the fluctuations.

According to our respondents, this benefit is real (Figure 12).

Figure 12

Roughly how would you assess the overall loading / utilisation level of your technology infrastructure (excluding people)?

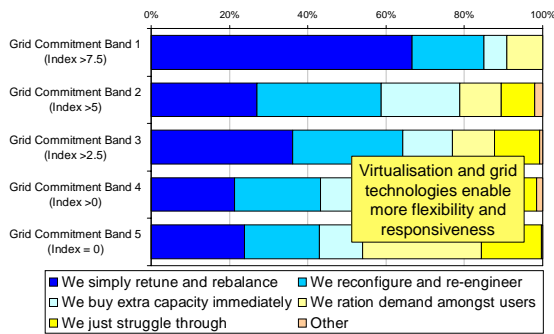


Unlocking more of the power available for useful work in this way means less hardware is required to deliver the same level of service, which in turn translates directly to cost savings.

Theory also says that the flexibility and configurability delivered by grid technologies means that the IT department can respond to growth and shrinkage in application requirements more effectively. There is again evidence of this from the research. When systems overload occurs, for example, those using grid related technologies tend to respond by simply retuning and rebalancing, whereas others are more inclined to take drastic measures such as immediately buying more hardware or even asking users to use the system less! (Figure 13).

Figure 13

When you experience overload problems, which expression best describes your approach to solving them?



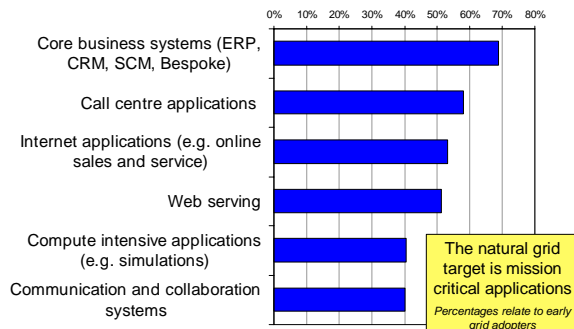
This is a clear example of how grid related technologies can improve the responsiveness of the IT department and their IT infrastructure to the needs of the business.

8. The Application of Grid

Given these benefits, it is not surprising that early grid computing adopters are targeting core business systems along with customer oriented applications such as call centre and online sales and service that require high performance and high availability (Figure 14).

Figure 14

Do you use Grid Computing in any of the following areas or are you likely to start doing so in the next 12 months?



This list also dispels the myth often propagated by the media that grid computing is only relevant to scientific applications. These would fall into the “compute intensive” category in Figure 14 and as we can see, this is fifth on the list of priorities.

9. Discussion

It is clear from this research that grid computing and its related component technologies in the virtualisation space are now on the mainstream agenda. Awareness is generally good and the more IT professionals get to know about these technologies, the more they appreciate the benefits.

We have also seen evidence that grid technologies, especially those concerned with virtualisation, are already being adopted to a significant degree. These trends suggest that many organisations around the world will be taking advantage of grid technologies sooner than some might have thought.

But there are some challenges that the industry and individual organisations need to address along the way.

Whilst the number of IT professionals getting the message has increased significantly, there are still concerns over the level of knowledge and skills that exist within organisations and the maturity of both the technologies and standards (Figure 15).

Figure 15

Which of the following would you regard as significant hurdles or inhibitors to investment in Virtualisation or Grid Computing technologies?



Progress in all of these areas is taking place rapidly, however, and the situation will settle over time, allowing even those averse to risk and change to take advantage of grid in the not too distant future. We have been through this cycle many times before with other emerging technologies and it is quite predictable, provided the benefits are confirmed, which they are in this case.

The last of our top 5 challenges is perhaps the most significant. As with any infrastructure related investment, there are sometimes difficulties in constructing a business case for grid related investments. This stems less from the problems of demonstrating a return on investment *per se* and more from the way in which budgets are structured and signed off in many organisations. The whole point of a grid computing environment is to achieve the smooth and dynamic sharing of resources between different applications and different areas of the business in response to fluctuating demand. Unfortunately, the consequence is that it becomes more difficult to determine who owns or should pay for a given component or operational framework, which is an issue when budgets are aligned to business units or departments.

This is one of the reasons why those further along with their adoption of grid highlight management of the political dimension as a challenge. In addition to the budget alignment issues there is also the natural association of grid computing with centralisation of IT resources and operations. This might not be obviously problematic in the early days, as grid technology is likely to be applied tactically in the first instance. As the threads are pulled together over time to create larger and larger grids (e.g. by joining smaller grids together) the resulting centralisation of activity potentially affects everything from internal cross charging mechanisms to reporting lines within management structures.

Whilst IT vendors and standards bodies get on with solving the technology related problems, organisations that are reliant on large and costly computing infrastructures would be well advised to start thinking about how they will deal with the softer issues arising from the kind of evolution we have been discussing.

One way or another, nearly all larger organisations will be impacted by grid technologies in the coming months and years. If its introduction is well managed the benefits will be manifest, as IT departments and the infrastructures they manage become better enabled to deliver value to users, whilst in the process becoming less and less visible.

APPENDIX A

A Business View of IT Infrastructure Evolution

The worse case IT infrastructure scenario, when viewed from a business perspective, is the existence of many different servers running a variety of equivalent but different applications in a completely uncoordinated manner. Whilst absolute chaos is thankfully relatively rare in the corporate environment, it is still common for organisations to be suffering from a significant degree of fragmentation. This is typically as a result of piecemeal evolution of IT systems in the data centre along with the legacy of decentralisation left over from the client/server revolution of the 1990s.

One of the main reasons we care about fragmentation from a business perspective is because it is inefficient, both in terms of hardware utilisation and overhead placed on the IT department. Fragmentation means a broader range of skills is required which translates directly to training and recruitment costs. It also means that similar tasks are duplicated for different systems – maintenance, support, upgrades and so on. Lastly, fragmentation invariably results in more interfaces between systems that need to be developed, maintained and managed from a risk of failure perspective.

Standardisation

The first step in dealing with such problems is to introduce a degree of standardisation. If one department or subsidiary is running Exchange 4.5, for example, another is running Exchange 2003, and yet another is running Lotus Domino, select a standard and move everyone to it. It is then possible to focus skill sets and rationalise support and maintenance processes. A similar standardisation process may be applied to line of business applications, such as ERP and CRM, and even platform components such as database management systems and application servers. In terms of the server hardware itself, it is possible to achieve a high degree of homogeneity, flexibility and manageability through the introduction of commodity based blade architectures. Standardisation in any of these areas will lead to better business resilience and lower overall operating costs.

Of course standardisation is easy to preach, but the associated switching costs and risks can often be significant. It is therefore important to realise that the immediate and direct benefits of standardisation are only part of the story. A better way, or perhaps equally valid way, of looking at standardisation is to view it as the first step towards the implementation of a genuinely efficient and flexible infrastructure.

Server Consolidation

Turning our attention to the platform level, i.e. the servers upon which applications are running, the problems of fragmentation and lack of standardisation are often extremely obvious. If you look around the average data centre, it will not be long before you come across a dedicated server that was originally installed to deal with the needs of a specific application for a specific department. Further investigation will then almost certainly uncover similar single function servers set up for other departments, possibly running a different version of the server operating system or a different brand of application, but essentially performing the same kind of task.

It is in this kind of extreme fragmentation situation that we typically find the poorest hardware utilisation. It is not uncommon for tens or even hundreds of servers in a data centre or spread around remote locations to be little more than ticking over most of the time, sometimes with utilisation levels as low as 5-10%. If these were employees who were only working productively for a fraction of the time we were paying them for, the organisation would be quick to rationalise, shedding excess capacity and keeping just enough employees to handle the workload comfortably but efficiently.

Such rationalisation is an obvious move in the context of the IT infrastructure. Modern server software can allow a single machine to be partitioned into multiple logical servers so that applications can coexist happily with other applications through giving each one its own “private” runtime environment. Standardisation at both a platform level (e.g. operating system, application server and database management system) and at an application level clearly opens up more consolidation opportunities as achieving efficient coexistence, or even consolidation at an application level (e.g. merging two previously separate accounting systems) is easier if there is a degree of consistency.

From a business perspective, the cost savings associated with server consolidation, both in terms of reduced capital investment and lower operational overhead, can be very compelling. Reports of an order of magnitude (10:1) reduction in the number of servers required to support the business have been reported by some who have gone down the consolidation route.

But the story does not stop with consolidation. This approach still ties an application to a discrete set of physical resources so the peak versus average sizing issue still lurks behind the scenes, albeit in a potentially diluted form. In order to break free of this constraint, we must look to an approach known as resource pooling and virtualisation.

Pooling and Virtualisation

A potential risk that arises when servers and applications are consolidated is the danger of having “all your eggs in one basket”. In the fragmented world, if a server went down, only one application in one part of the business would be affected. If we have gone through an aggressive consolidation exercise, the likelihood is that business impact of a server failure will be significantly greater.

Fault tolerance is therefore a hot topic for organisations going through the consolidation process.

One way of tackling this problem is to look to the “top of the range” highly fault-tolerant servers offered by mainstream hardware vendors. However you measure it – mean time between failure, system availability, etc – these systems are extremely resilient so can provide a good degree of business confidence.

They can also potentially better deal with the more concentrated processing and storage requirements that are the natural result of consolidation activity as they pack a lot of horsepower into a single box.

Ironically, however, building systems that take fault tolerance and scalability to the next level again is probably best achieved by going to the other extreme completely. Some of the most compelling developments in the IT industry have been concerned with allowing many smaller machines to work together in a coordinated manner to deliver a pool of processing power that may be viewed as a single resource entity. To put it another way, to make a lot of small commodity computers look like one big virtual server.

The key to implementing this kind of architecture is systems software that is able to use knowledge of the resources available in the resource "pool", monitor the loading on each component and distribute work in a balanced manner. Such software also ensures that when a component fails, the work of that component is reallocated seamlessly to other components, which achieves the required levels of fault tolerance.

This virtualisation capability has been incorporated into solutions from Oracle, Microsoft, IBM, SAP, HP, CA, Citrix and other mainstream IT vendors, as well as a number of specialist niche players. Sometimes it is positioned as an explicit offering, sometimes as a feature of broader solutions.

Storage vendors too, such as Symantec/Veritas, EMC and StorageTek, have embraced the concept of virtualisation, allowing multiple physical storage devices to be managed as one or more virtual devices.

Whatever the context, the net effect of virtualisation is to sever the traditional hardwired link between applications and the individual underlying physical resources that provide the computing power and data storage capacity they need.

Apart from enhancing business resilience and reducing costs yet again by further simplifying the application deployment, management and maintenance process, virtualisation also means that the IT department may be much more responsive to the needs of the business.

Virtualisation allows IT personnel to allocate or de-allocate physical resources to or from a particular application much more easily. No longer is the lengthy and labour intensive process of configuring the application to recognise and make use of the new component required. Neither is the work associated with freeing up a component that is no longer needed without disrupting the application. New physical resources, whether they are servers or storage components, are simply added or removed from the pool through processes known as "provisioning" and "deprovisioning" respectively. Meanwhile, the application continues running in blissful ignorance of what has gone on behind the scenes.

The spin-off benefit of virtualisation is that organisations may use inexpensive commodity components to achieve the same level of scalability and fault tolerance offered by high end servers, but at an overall lower cost.

Grid Computing

Virtualisation techniques are useful to manage the steady growth and contraction of application related activity as macro usage patterns change. But in itself, it does not deal with the more granular fluctuations in demand that take place across all application areas over the course of the working day or working week, for example.

We have already discussed the concept of peaks in activity, such as the intense usage of email at the start of the day. There are many examples of similar peaks in other areas, e.g. the consumer focused call centre system that is under

pressure in the middle of the day and early evenings as customers make contact during their lunch break or when they get home from work. Demand on financial applications can similarly be cyclic, with different activities causing peaks on a daily, weekly, monthly or quarterly basis.

In addition, we also have one-off requirements to take account of. An example of this might be a product launch, for which extra contract telemarketers are hired to drive an intensive two week outbound sales effort. This could place more demand on call centre systems requiring extra computing power for the duration of the exercise that may be released once it is over.

Dealing with cyclic or transient fluctuations in demand that take place over a short timescale is where grid computing becomes relevant. The idea is to take virtualisation to the next level by adding a degree of automation. In its ultimate form, grid computing allows an organisation to "pool the pools" of physical computing resource and allow the allocation and deallocation of components to take place dynamically across all application areas.

It is here that we need to switch analogies from the pool concept to the grid concept. The term "grid" originates from the power industry. In that context, we have multiple power stations generating electricity and contributing power to the electricity grid. For a consumer to make use of that electricity, they simply plug their appliance into a socket connected to the grid and get the power they need, with absolutely no knowledge of where that power comes from or how it is delivered. They do not care about how the output of particular stations is blended or managed to deal with peaks and troughs of overall demand. Neither do they care about how the grid is managed to ensure that an individual power station outage is taken care of by spreading the demand it would normally deal with across other stations.

With grid computing, the application is equivalent to the consumer in our power analogy and in the same way, it metaphorically "plugs itself" into the grid to obtain the computing power and storage capacity it needs. Behind the scenes, work is dynamically allocated to physical resources as the demand profile changes. An individual server could, for example, be supporting email between 9am and 6pm but be providing additional power to the call centre management system to deal with the peak there between 6pm and 8pm when the email load is much lighter.

If the grid computing concept is new to you, this can sound a bit far-fetched. Without going into too much detail, we will therefore take a brief look at the underlying mechanics.

Grid computing relies on the ability to rapidly provision and deprovision servers (or server partitions) to run a particular application. Sometimes this may require the whole software stack to be installed onto the machine – operating system, application server and application, for example. Depending on what the server was doing before, however, it may just be necessary to load the application from shared storage and run it. Figuring out these dependencies is the job of the grid management and control software that enables the grid computing approach. The upshot is that a manual installation process typically taking hours to complete is replaced by an "intelligent" automated process that executes in minutes (some vendors are even talking about seconds to provision a server).

The grid management and control software also allows us to monitor the resource demands from applications and the loading on individual components that make up the grid. This

theoretically gives it the ability to move resources around fully automatically in a near-real time manner to ensure that the power is delivered where it is needed rather than sitting idle in a less active area. In practice, systems administrators may want to maintain a degree of manual control over the resource allocation process, but even so, automatic provisioning means this may be achieved by simply moving resources around via a console rather than having to physically deal with the details of software installation and server configuration. In addition, the administrator's life is made easier through the grid management and control

software raising alerts when imbalances arise and then making recommendations or providing options for re-balancing the environment.

From a business point of view, grid computing is a way of ensuring that all of an organisation's computing assets are used together to maximum effect, delivering the ultimate combination of fault tolerance, cost efficiency and service levels to the business.

APPENDIX B

Origin of the Grid Indices**Approach and Rationale**

Assessing and tracking the level of progress and adoption of emerging ideas and technologies is a challenge. Measuring the end point only, e.g. the number of companies achieving return on investment, can be potentially misleading. Using such an approach, it would be difficult, for example, to distinguish between a technology with broad appeal that was at the beginnings of widespread adoption, versus a niche technology that had achieved rapid adoption in certain areas but had limited growth potential beyond this.

It is for this kind of reason that the overall Grid Index is based on more detailed measurements based on activity throughout the adoption lifecycle – the acquisition of knowledge (knowledge index), appreciation and/or achievement of benefits (benefit index), activity or commitment to act (commitment index) and achievement or expectations of Return on Investment (ROI index).

By correlating these, we can both validate our assumptions of the lifecycle (e.g. those who understand tend to move on to appreciate the benefits, those who appreciate the benefits tend to act soon after, etc) and track how activity is unfolding long before the end point is reached.

The absolute values of the indices themselves are far less important than the differences between them and how they change over time. The Grid Index programme is therefore designed to run on a regular basis and as time goes on, we will be able to “calibrate” the indices more accurately and continue to develop benchmarks by looking at scores achieved by the most committed adopters. We can learn a great deal from comparing and contrasting index values as we have seen in this report.

Data Gathering and Index Calculation

Questions used to capture information during telephone interviews were designed to allow the respondent to provide a meaningful response, but also allow that response to be translated into a numeric score, e.g.

| How would you rate your level of knowledge and understanding of the concepts and practicalities of Grid Computing? | |
|---|------------|
| Deep understanding | Score = 10 |
| Comfortably familiar | Score = 8 |
| Limited familiarity | Score = 6 |
| Passing awareness | Score = 4 |
| Negligible knowledge | Score = 2 |
| Completely unfamiliar | Score = 0 |

Using this kind of technique, similarly formatted questions were constructed in the following areas:

Knowledge Index

- Level of knowledge and understanding

Benefit Index:

- Level of benefit in relation to better operational efficiency
- Level of benefit in relation to improved service to users

Commitment index:

- Current status of adoption
- Part played in overall IT strategy
- Likely investment over the next 12 months

ROI index

- Timeframe for achieving tangible return on investment

These questions were asked separately for Server Virtualisation, Storage Virtualisation and Grid Computing. The results were then aggregated and averaged to produce the higher level indices.

The approach to deriving indices for Standardisation and Consolidation was very similar, apart from the technologies that were highlighted as part of the question. Rather than Virtualisation and Grid technologies, questions focused on traditional areas such as application servers, database management systems, email environments, etc.

APPENDIX C

Interview Sample Distribution

The information presented in this report was derived from 1356 interviews with senior IT influencers and decision makers completed in March 2005. 300 of these were conducted in North America, 600 in Europe and the remainder (456) in Asia Pacific. Respondents were from a mixture of large multinationals and medium to large national organisations across a broad cross section of industry sectors.

Distribution of the sample by geography and industry was as follows (Figures 16 and 17):

Figure 16

Distribution by Geography





Figure 17

Distribution by Industry



About Oracle

Oracle (NASDAQ: ORCL) is the world's largest enterprise software company. For more information about Oracle, please visit <http://www.oracle.com>.

ORACLE®

About Quocirca

Quocirca is a research and analysis company with a focus on the European market for information technology and communications (ITC). Its analyst team is made up of real-world practitioners with first hand experience of ITC delivery who continuously research and track the industry in the following key areas:

- Business Process Evolution and Enablement
- Enterprise Applications and Integration
- Communications, Collaboration and Mobility
- Infrastructure and IT Systems Management
- Utility Computing and Delivery of IT as a Service
- IT Delivery Channels and Practices
- IT Investment Activity, Behaviour and Planning

Quocirca research is always pragmatic, business orientated and conducted in the context of the bigger picture. ITC has the ability to transform businesses and the processes that drive them, but often fails to do so. Quocirca's mission is to help its customers improve their success rate.

Quocirca has a pro-active primary research programme, regularly polling users, purchasers and resellers of ITC products and services on the issues of the day. Over time, Quocirca has built a picture of long term investment trends, providing invaluable information for the whole of the ITC community.

Quocirca works with global and local providers of ITC products and services to help them deliver on the promise that ITC holds for business. Quocirca's clients include Morgan Stanley, Oracle, Microsoft, IBM, CA and Cisco. Sponsorship of specific studies by such organisations allows much of Quocirca's research to be placed into the public domain. Quocirca's independent culture and the real-world experience of Quocirca's analysts, however, ensures that our research and analysis is always objective, accurate, actionable and challenging.

Many Quocirca reports are freely available and may be requested via www.quocirca.com.

Contact:

Quocirca Ltd
Mountbatten House
Fairacres
Windsor
Berkshire
SL4 4LE
United Kingdom

Tel +44 1753 754 838
Email info@quocirca.com

The logo for Quocirca, featuring the word "quocirca" in a lowercase, sans-serif font. The letters "qu" are in blue, "o" is black, "c" is red, "i" is black, "r" is red, and "ca" are black.