SPIRENT

Issue 2

Navigating the Path to Optimal Cloud Services

Featuring research from

Gartner

Navigating the Path to Optimal Cloud Services

Provisioning, configuration, regulatory compliance, incident, change and problem management consume a huge portion of a typical IT budget and are traditionally administered manually in silos with proprietary, incompatible tools.

In recent years, the concept of cloud computing has steadily moved to the forefront of concern for IT professionals. The field is an amalgam of processes, technologies, platforms and sometimes confusing or imprecise terminology, all with common goals. These goals include: reducing the labor-intensive, tactical aspects of providing IT services, optimizing teams and resources, maximizing hardware investments and amplifying speed of services delivery. As a result managers are freed to pursue innovation and strategic planning.

The dynamic data center incorporates one or more elements, including:

- Public or private cloud services - self-service ordering, real-time provisioning, and metering
- IT service management - automated and proactive monitoring, analysis and management of service performance, availability and scalability.
- Real-time infrastructure (RTI) - runtime optimization and execution of resources at the service and subservice levels

RTI is a basic building block of the dynamic data center, the fundamental level of IT element automation. It maps the demand for shared services and resources to the supply, making possible the elastic provisioning and scaling of resources for services to meet SLAs efficiently and economically.

A review of past and recent research reveals that while many offerings are at the Technology Trigger point of the Gartner Hype Cycle[1] (the first of five stages in the cycle of IT innovation), the trend indicates a significant level of adoption of RTI as some organizations have already moved from discussion to implementation.

Spirent believes that the Gartner survey research that follows highlights the importance of network assessment during planning, development, deployment and ongoing operation to identify issues that affect scalability, interoperability, availability and performance.

To help illuminate the most productive path to implementation of cloud services, Spirent Communications delivers knowledgeable perspectives on the directions early cloud adopters are taking in this publication featuring Gartner research. The Gartner research provides recommendations regarding foundations for cloud development and strategic first steps. RTI and cloud services hold the promise of maximum scale with minimum administration. These findings can guide you around pitfalls and smooth the way to achieving those benefits. Spirent also provides guidelines for selecting the best qualified partners for a successful implementation.

Source: Spirent

FIGURE 1

Three Different Disciplines Coming Together to Optimize Runtime Services: IT Service Management, Cloud Services and RTI

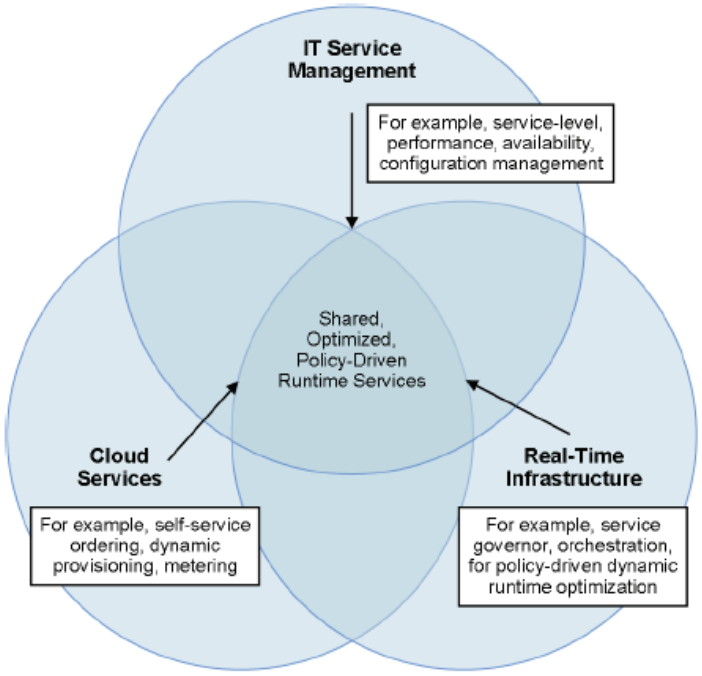

Source: Gartner (March 2010)

[1]Gartner RAS Core Research Note G00174718 Building Private Clouds With Real-Time Infrastructure Architectures, Donna Scott, 22 March 2010

Survey Shows High Interest in RTI, Private Cloud

During the presentation "Building the Private Cloud With Real-Time Infrastructure Architectures," we surveyed the audience on use cases, drivers and inhibitors. We analyze the results in this research.

Key Findings

- 24% of the audience (13 out of 55 responses) have already implemented private clouds for development and test labs; more than double that number plan to implement private clouds in the next two years.
- Real-time infrastructure (RTI) has advanced in recent years, with 26% of the audience already having implemented RTI architectures.
- Attributes of public cloud computing, such as self-service ordering and dynamic provisioning, are also sought internally for private cloud services to speed service delivery and reduce costs. 36% of the audience cited speeding service delivery as their No. 1 driver for private cloud computing.
- The primary use case for RTI is infrastructure optimization through virtual server mobility.

Recommendations

- Pilot private cloud and RTI implementation in less-mission-critical lab and test environments first.
- Implement metrics, including cost, mean time to change and mean time to provision, in order to remove cultural biases with numbers and facts.
- Change the asset funding model from new-project-based funding to one that is focused on buying in anticipation of demand, based on forecasting and demand management.
- Get commitment for funding to be made as capacity is needed, rather than at the time of implementation.
- Use the Gartner Infrastructure and Operations (I&O) Maturity Model to be sure you are at Level 3 (proactive), prior to implementing private cloud services or RTI architectures.

STRATEGIC PLANNING ASSUMPTION(S)

Through 2014, data-center-wide RTI will not exist.

ANALYSIS

In interactive electronic polling surveys conducted at Gartner's December 2009 Data Center Conference in Las Vegas, attendees consisting mainly of IT operations and data center managers and other senior managers were asked how they see RTI and private cloud environments evolving. The sample sizes varied from 40 to 60 individual responses and are not statistically extensive distributions, but the results are of interest to anyone pursuing IT operations management and runtime optimization strategies. Furthermore, these results are consistent with Gartner's client inquiries and research. We used weighted averages where the audience had multiple choices.

Status With Private Clouds and RTI

The first two questions focus on penetration of private clouds and RTI architectures within the audience's data centers. RTI provides runtime optimization and execution for services and resources. RTI maps the demand for shared services and resources to the supply of shared resources, enabling a more efficient and cost-effective runtime environment. Private cloud services are often targeted toward high-volume service requests; their self-service and dynamic-provisioning attributes make them well-suited to this type of service model.

For more information on the similarities and differences between private cloud and RTI, see "Building Private Clouds With Real-Time Infrastructure Architectures."

As shown in survey results, given the nascent nature of much of the technology in this space, the audience was fairly sophisticated, with 13% already having implemented private clouds in production and 26% having already implemented RTI architectures. In Figure 1, because the question allowed more than one reply, a total of 68 responses were received, with 30 of those indicating implementation of private clouds overall (or 44%), primarily in nonproduction use cases. Within private cloud implementations, just 23% of attendees (or seven responses out of 30) were for production purposes, such as implementation of self-service request and provisioning for Web or database services. 77% were for nonproduction services, including development/test labs or demonstration/training labs. It's typical to trial emerging technologies first in nonproduction environments prior to production implementations. Many IT organizations, however, have set their sights on private cloud implementations, with more than 50% of attendees (29 of 52 responses) planning to implement during the next two years. This shows how desirable cloud attributes are to IT organizations – with self-service ordering and dynamic provisioning of standardized catalogued services both increasing IT agility and reducing labor and capital costs through standardization and automation.

With regard to RTI architectures in Figure 2 (where only one response was allowed), of the 26% of the audience that had already implemented them, there was an even split between production and nonproduction use cases, with a small percentage (2% of the audience) focused on repurposing disaster recovery (DR) environments. Clearly, the audience responded on the development/test use case for both private cloud and RTI – with both represented at 24% (13 out of 55 responses) and 12% of the audience,

FIGURE 1

What is the status of your private cloud initiatives?

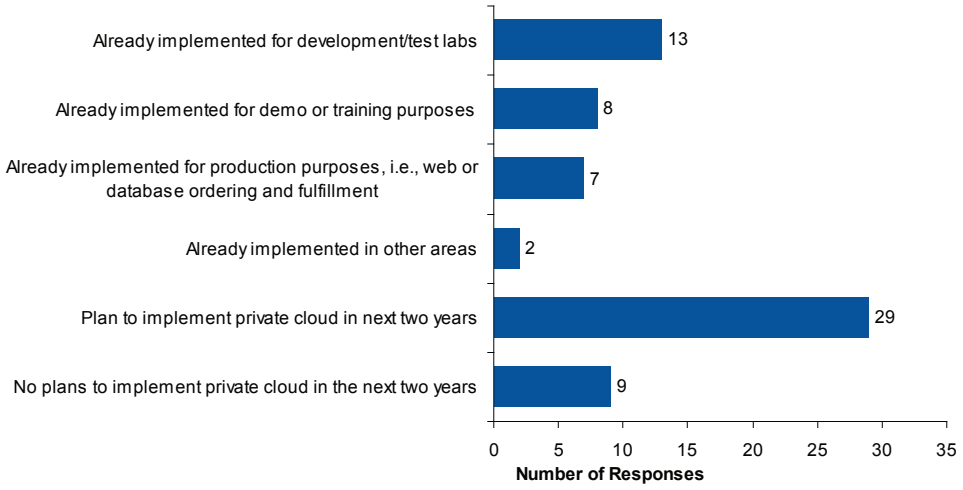

The audience could vote for more than one; therefore; the actual votes exceed the number of interactive polling devices, which were 55. (DB = database)
Source: Gartner (March 2010)

FIGURE 2

Have you implemented RTI?

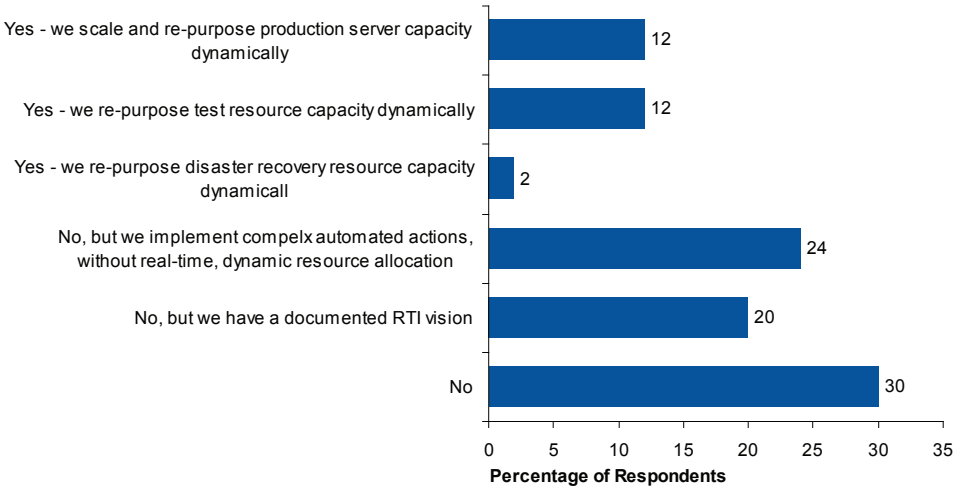

Source: Gartner (March 2010)

respectively. 24% of the audience implements complex automation without dynamic resource allocation, most likely through run book automation and orchestration, which often is a precursor to RTI. Moreover, 20% of the audience has a documented vision for RTI. 30%, however, have not planned for RTI in their data centers during the next two to three years. This group may be technology followers, waiting for others to pave the way to value first; or they may be investing in the prerequisites in the area of process maturity necessary to succeed with RTI (see the IT I&O Maturity Levels section). Another interpretation could be that the value has not been proved; there is insufficient data to provide a conclusive interpretation.

Drivers for Private Cloud Implementations

Two key attributes of public cloud computing are self-service ordering and dynamic provisioning – these same attributes can enable the internal IT organization to increase agility and speed of service delivery when implementing private cloud computing. As shown in Figure 3, this was the main driver for the audience. For example, many development/test organizations complain that it takes two or more months to provision the operating environments necessary for them to be efficient and effective in their work. Implementing a private cloud service can reduce this time frame to a matter of hours, and can give users more control over planning and requesting resources (which is why 24% of respondents have already implemented private cloud for development/test and many more plan to do so). Moreover, costs are reduced not only due to shared resources (with allocations based on policy), but also because labor is taken out of the service delivery process. Cost was cited as the second driver for moving to private clouds. Those citing increased quality and consistency likely want to gain this benefit from the standardization enabled and enforced in a private cloud implementation.

FIGURE 3

What is your main driver for moving to private clouds?

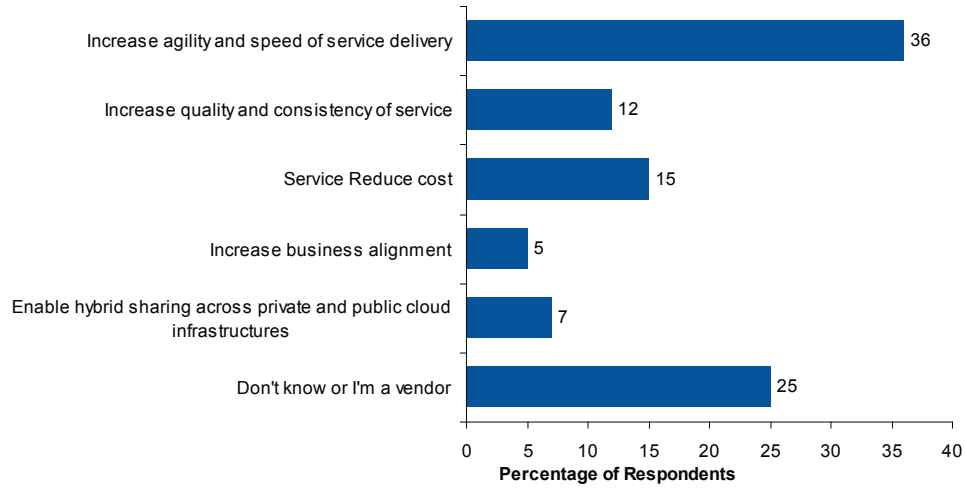

Source: Gartner (March 2010)

FIGURE 4

What are your focus areas for RTI?

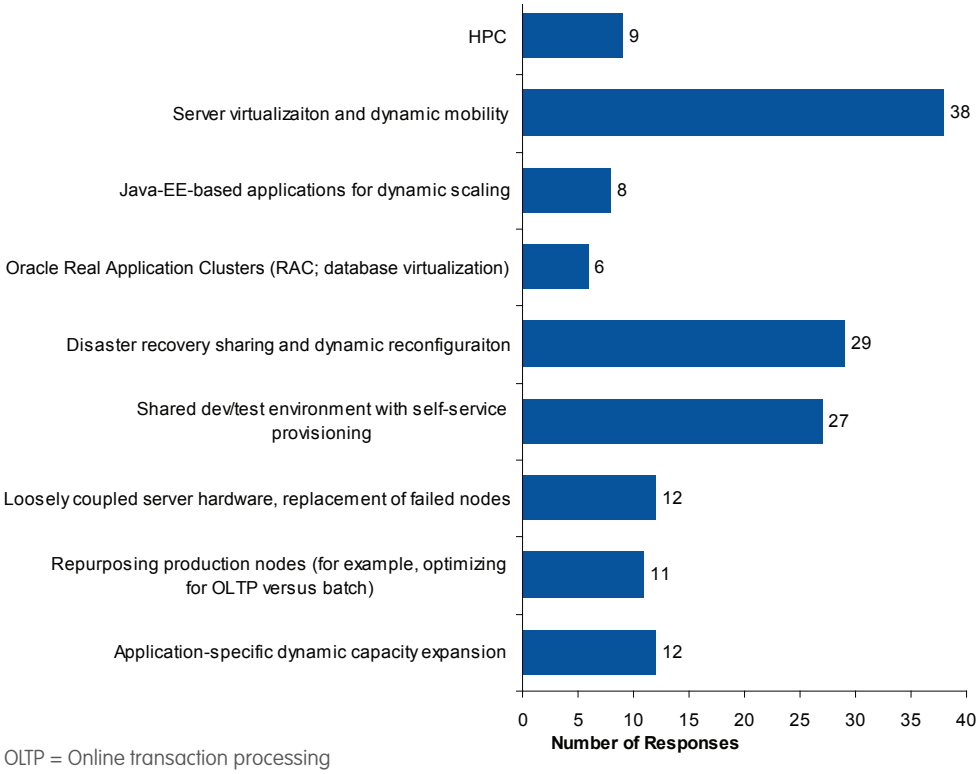

OLTP = Online transaction processing
Source: Gartner (March 2010)

Use Cases for RTI

RTI focuses on optimizing the runtime environment by enabling better resource utilization to services dynamically to meet SLAs and business priorities. As such, RTI can be implemented in a shared hosting environment even for custom, one-off services or applications being hosted in order to share capacity with each other and increase asset utilization. Figure 4 depicts the audience's use cases for RTI, either implemented or planned. 78% are focusing on using server virtualization mobility (such as VMware's VMotion or Microsoft's Live Migration) to drive greater infrastructure optimization. This is most likely due to the increased presence of server virtualization, as well as the audience's area of responsibility, which is most likely I&O. While we agree that this use case can and should be leveraged, we also recommend that a more service-centric focus be implemented to optimize availability and performance.

The second most cited RTI use case is DR reconfiguration, at 59% (29 out of 49 total responses). DR continues to be a priority for many organizations, especially with shortening recovery time objectives. Using RTI can enable a secondary data center to share development/test environments with DR so that in the event of a disaster or for a DR test, servers are reconfigured to look like the production environment. Critical technologies supporting this capability are server provisioning and configuration management tools and repurposing technologies (which are based on imaging and virtual infrastructures).

Following right after in priority as an RTI use case is self-service development/test lab provisioning, at 55% (27 out of 49), which were highlighted in the prior two questions as well. All other responses received less than 25% of the audience vote. However, it is important to note that the low counts

could be due to the audience responsibility areas, which may not include grid/high-performance computing (HPC), Java Platform, Enterprise Edition (Java EE) and other types of applications that may benefit from RTI.

IT I&O Maturity Levels

Gartner's I&O Maturity Model can be used as a readiness assessment for private cloud and RTI. We recommend that organizations be at a proactive level of maturity (Level 3) or above to be successful with private cloud and RTI. This is because critical processes such as change and configuration management are mastered at Level 3, and without them private clouds and RTI would likely fail in a production environment. Organizations that cannot manage change in a manual way now will have even more difficulty when more dynamic changes are occurring. The exceptions are for nonproduction environments (such as development/test labs) where less rigor is required.

As shown in Figure 5, 58% of this audience ranked their organizations at Level 3 or higher. Either the audience was more advanced (based on the nature of the presentation they chose to attend), or they are overestimating their maturity levels, Gartner's more-detailed online assessments of I&O maturity show only 11% of organizations at proactive maturity and above. Due to the high levels of private cloud and RTI implementations, we believe the reason for this audience's ranking of their organizations is that they are more advanced than average.

Inhibitors to RTI and Private Cloud

We asked about hurdles in implementing private cloud and RTI architectures. The results are shown in Figure 6; since the audience could choose up to three hurdles, the results are "weighted," thus the "w" means weighted results. Interestingly, the top four results are all business management and organizational issues, and are not

FIGURE 5

What is your estimated of I&O maturity level?

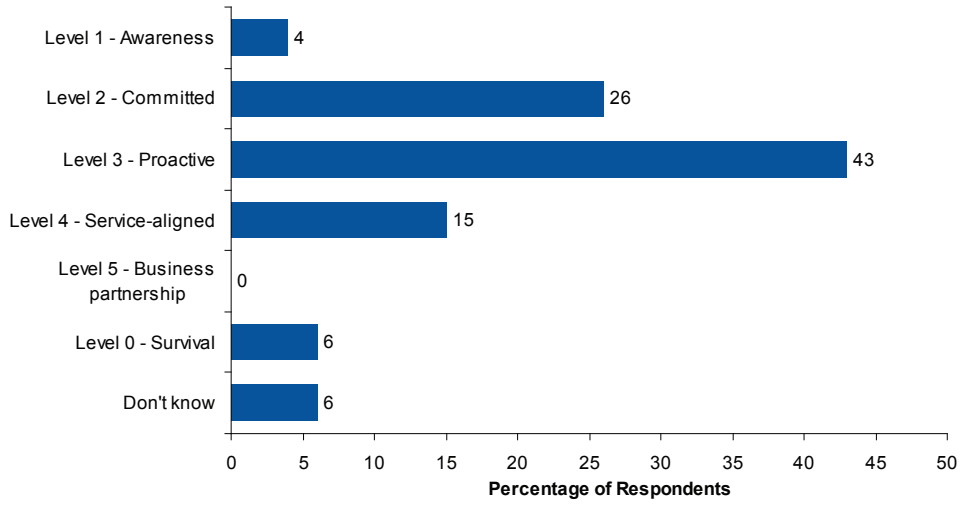

Source: Gartner (March 2010)

FIGURE 6

What are your primary inhibitors to RTI and private cloud?

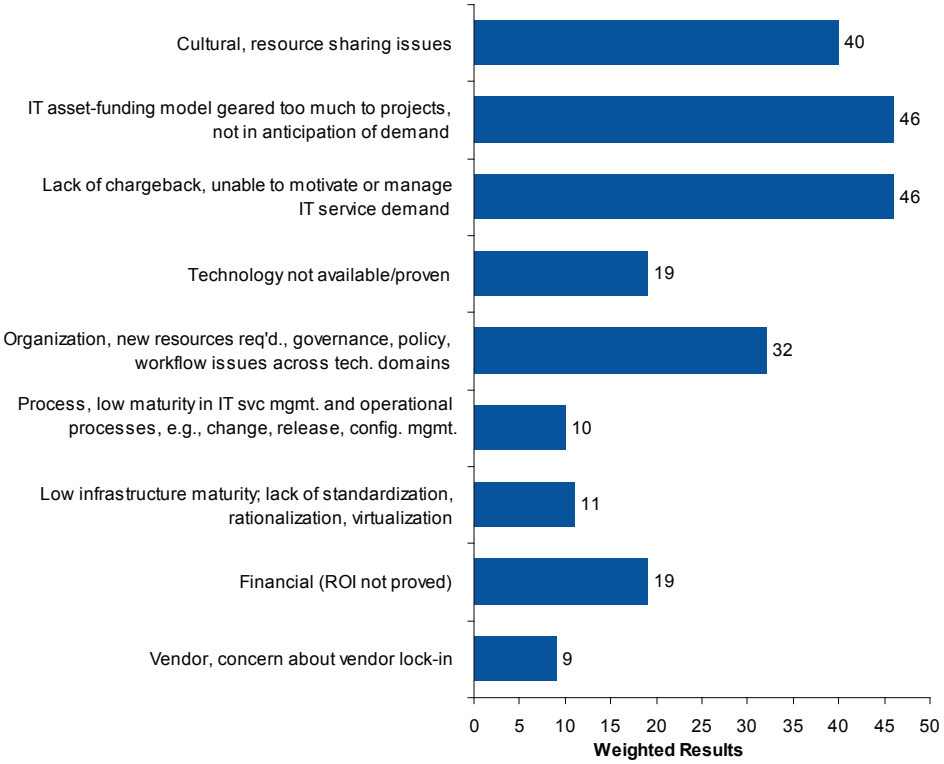

ROI = Return on investment
Source: Gartner (March 2010)

related to technology issues. Tied at No. 1 are the IT asset funding model and lack of chargeback. Many IT organizations do not have any chargeback mechanisms, thus, they tie capital asset investment to projects, rather than to capacity and demand management processes. As a result, shared consolidated resource pools may be implemented, but the sponsoring organization does not have the means to increase the size of the pool of resources as demand grows. Ideally, the funding model changes to one of forecasting and demand management (as opposed to solely project-based). However, even project-based funding can get around some of these issues by getting approval to buy capacity for forecast demand (either at project implementation time or staged as needed) – but every service in the shared pool must be analyzed and assessed, or new projects will wind up paying for the increased capacity of existing services, which would discourage the use of the shared environment.

Closely related is lack of chargeback – or worse – lack of knowledge of the cost of service delivery. Business decisions include cost and benefit data; however, data center hosting costs often are not known. As a result, hosting is perceived as "free" by customers, which results in the inability to manage demand and to effectively plan for facility and resource capacity. Charging back, or at least showing customers what the costs are, gives them the ability to better assess their needs and priorities, and make more effective business decisions. It also enables more-accurate comparisons of hosting costs between internal and external service providers.

The third inhibitor cited relates to culture, specifically customers not wanting to share resources with each other. This sometimes derives specifically from a project-based funding model ("I paid for the resources and don't want to share them with others"). To overcome resistance, we recommend setting up a pilot environment (for example, a test/development private cloud implementation) to prove how it can save money, speed service delivery and increase quality of service.

The fourth inhibitor is organizational. Virtualizing and sharing resources requires that technology domains work together. For example, to provision servers quickly requires not just standardization of the software stack, but also that the server, network and storage team work together to integrate their processes and policies. This is rarely done well across technical domains and organizational silos – but for virtual environments it is even more critical.

Vendor Confidence

Our final question asked which vendors the audience had confidence in to deliver on private cloud computing and RTI. While the sample size is low, the results show that nearly 50% of the audience relies on VMware to achieve their private cloud and RTI objectives. This makes sense when you consider that the primary focus for RTI is virtual server mobility, and many in the audience may also be implementing VMware for test/development lab management. In addition, virtualization is often a precursor to cloud implementations. No other vendor achieved a significant percentage of responses, although the collection of infrastructure vendors – Cisco-EMC, HP, IBM, Microsoft, Novell and Oracle —add up to 39%. This may suggest that some enterprises are counting on their key hardware and software infrastructure providers to provide optimization for their data centers.

When infrastructure vendors provide such functionality, they typically offer it for their own infrastructures, or at least to optimize the best for their own infrastructures. Finally, BMC Software's and CA's low rankings may reflect the limited visibility and perception customers have with respect to their work in virtualization. In addition, management software vendors do not have mind share in runtime execution outside of job scheduling.

Bottom Line

Server virtualization is becoming more common, and, as the results show, enterprises are staring to prepare and invest in getting more out of it than just the benefits of consolidation. They are seeking greater agility levels and lower labor costs through the implementation of private cloud services and RTI architectures. Survey data shows an advanced audience, with 30% having already implemented private cloud services and 26% having already implemented RTI architectures. We recommend that enterprises seeking to implement private cloud and RTI assess themselves against the Gartner I&O Maturity Model to be sure they are at Level 3 (or higher) to have the best chance of success.

Gartner RAS Core Research G00174889, Donna Scott, 26 March 2010

Building World-Class Cloud Services: Essential Steps

Virtualization Renders Legacy Test Methodologies Incomplete

The preceding Gartner survey results revealed that the primary use case for RTI is infrastructure optimization through virtual server mobility. It also indicated that optimization is achieved through an iterative development process of configuration and assessment.[2] However, legacy test methodologies are based on physical testers connected to physical ports on the device or system under test. Virtualization involves multiple virtual machines (VMs) that may reside anywhere in the cloud, on separate physical servers or on a single physical server.

When processes involve VMs on the same physical server, a physical tester (dependent on a physical port for visibility into the system) cannot assess intra-server performance, availability, security, or scalability. When assessing a pilot system in a less-mission-critical lab or test environment (as recommended by the survey), testers must be at the endpoints to generate traffic and accumulate results.

So even in the case of *inter*-server assessment of virtualization, when the endpoint is a VM, a physical tester is irretrievably relegated to a physical port connection. This position of the tester is not at an endpoint but in the middle, and it is more suited for physical infrastructure testing as opposed to virtual infrastructure testing which could include:

- Virtual firewalls
- Virtual switches
- Virtual proxies
- Virtual load balancers

End-to-end assessment of a virtualization implementation requires a virtual tester inside the VMs. A virtual tester is a software-based test system implemented in a virtual machine. To the network devices under test, and to the test engineer, it looks and behaves exactly as if it were a hardware tester. A virtual tester makes it possible to test virtualization at all the levels it has impact: intra-server, inter-server and infrastructure.

Comprehensive Testing Assures Mature Technology

Delivering cloud services involves a complex and possibly confounding convergence of processes, technologies, platforms and devices, many of which are in the early stages of adoption. Of course, all systems must be assessed and validated, even those built on mature technologies, and even more so systems that incorporate nascent technology.

Existing technologies reached maturity through the development of successive test methodologies that assessed and validated specific implementations. The evolving elements of cloud services will mature only through rigorous testing with new methodologies designed specifically to reveal their weaknesses, and strengths. According to a recent Gartner study on Cloud Computing, server-based computing solutions are very sensitive to network performance. The study also 'recommends being cautious with vendor claims, particularly regarding scalability, availability and security.[3]

Recognizing and responding to these vulnerabilities, Spirent Communications developed the industry's first holistic test methodology to validate the performance, availability, security and scalability (PASS) of cloud computing. The Spirent PASS methodology includes both physical and virtual appliances specifically designed to test services and infrastructure between any points in the cloud environment. Details of the PASS methodology include:

- **Performance:** Optimize cloud services and infrastructure to maximize user experience
- **Availability:** Ensure high availability in daily operation and under disaster conditions
- **Security:** Eliminate vulnerability and exposure between tenants in the cloud
- **Scalability:** Validate responsiveness as demand varies according to tenant needs

Gartner also recommends that a more service-centric focus be implemented to optimize availability and performance.[4]

PASS incorporates a wide range of use cases to assess the application components and network elements of a private, public or hybrid cloud. PASS delivers unprecedented granularity through use cases that subject the cloud service to inspection at the infrastructure level (end to end), the inter-server level (applications, services, VMs on different physical servers), and intra-server level (between VMs within a physical server).

For example, at the service level, employing the PASS model assesses user quality of experience (QoE) under realistic and peak subscriber loads, scaling to multi-tenant loads, under the stress of failover scenarios, while subjected to threats mixed with encrypted traffic.

At the infrastructure level, PASS assesses end-to-end quality of service (QoS) at critical points:

[2]Gartner RAS Core Research Note G00174889 Survey Shows High Interest in RTI, Private Cloud, Donna Scott, 26 March 2010
[3]Gartner RAS Core Research Note G00174116 Understanding 'Client in the Cloud', Stephen Kleynhans 10 May 2010
[4]Gartner RAS Core Research Note G00174889 Survey Shows High Interest in RTI, Private Cloud, Donna Scott, 26 March 2010

FIGURE 2
Unknowns of Clouds Computing

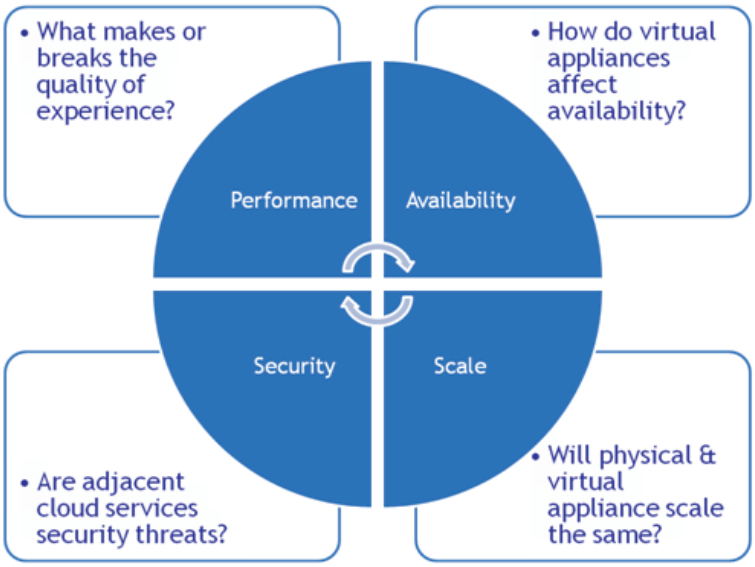

Source: Spirent (May 2010)

- **Network elements:** high-density 10/40/100 Gb Ethernet, converged FC and Fiber Channel over Ethernet (FCoE)
- **Application elements:** firewalls, intrusion prevention systems, WAN accelerators, proxy servers, and others
- **Virtual elements:** virtual switches, virtual firewalls, virtual load balancers, and virtual proxies, via a combination of virtual and physical test systems

Spirent's development of the PASS test methodology is a milestone in the journey from innovation to adoption of cloud services technologies, and a dynamic facilitator for those looking to establish world-class cloud services.

Assessment of Cloud Services Is Essential at Multiple Points

The Gartner survey identified multiple points where assessment of the performance and availability of a cloud services implementation assures successful development and deployment. Other Gartner findings on the value of testing include:

- **Gartner Analysis:** *Every service in the shared pool must be analyzed and assessed, or new projects will wind up paying for the increased capacity of existing services, which would discourage the use of the shared environment.*[5] **Spirent insight:** Assessment of service utilization establishes a baseline for every component in the system and documents per-service utilization for equitable funding decisions.
- **Gartner analysis:** *No other vendor achieved a significant percentage of survey responses, although the collection of infrastructure vendors . . . add up to 39%. This may suggest that some enterprises are counting on their key hardware and software infrastructure providers to provide optimization for their data centers.* **Spirent insight:** There are a number of vendors in the pack, but no clear leaders in the field. Testing identifies the right fit for each cloud. In addition, assessment is best performed by a neutral third party, not a vendor with vested interests.
- **Gartner Recommendation:** *Get commitment for funding as capacity is needed, rather than at the time of implementation.*[6] **Spirent insight:** Having approved funding for timely performance and scalability testing enables managers to identify the breaking points of the system and quantify the delta between current utilization levels and maximums, avoiding any costly and debilitating problems after launch. Testing also indicates the headroom in the current implementation and deployment.
- **Gartner Recommendation:** *Pilot private cloud and RTI implementation in less-mission-critical lab and test environments first.*[7] **Spirent insight:** To achieve this pilot initiative, holistic testing is one of the primary components of success. The test bed is the ideal environment to develop and optimize both test procedures and service implementation.

[5]ibid
[6]ibid
[7]ibid

Planning the Implementation of Cloud Services

Implementation of IT initiatives such as RTI, virtualization and cloud services represent significant infrastructure and budget commitments with considerable risk. According to Gartner, a common response to development slips is to compress the testing schedule. Overall, downplaying the roll and shortchanging the schedule ends up costing organizations considerably more in the long term, with increased maintenance and defect fix costs.[8]

Best practices clearly indicate testing to mitigate these risks. However, relying on in-house engineering resources that lack the required expertise often produces unreliable test results. The choice of a qualified test partner is critical. Employing a partner with a team of test experts conversant in the challenges of cloud computing with expertise in the field of data center testing is essential. Also, since the cost of the required test equipment can be prohibitive, selecting a partner that also manufactures a broad array of networking test equipment, and provides rental options, offers added advantages and cost benefits. Spirent Communications with their world-wide Professional Services group is an example of a third-party partner with all these qualifications, as well as those listed below.

Additional attributes to look for when choosing a testing partner include:

- Testing as a core competency
- Virtual and physical test systems for true end-to-end testing of cloud services
- An established name in the test and measurement industry
- A neutral third party
- Experience and expertise
- Quality test engineers
- An established delivery process
- Extensive automation expertise
- Articulation of testing benefits / ROI vs Risk
- Holistic understanding of networks
- Successful engagements with references

Conclusion

When implementing private cloud and RTI architectures, Gartner recognizes that every service in the shared pool must be analyzed and assessed.[9] Spirent's experience indicates that a neutral, qualified third-party test solution is recommended for the successful implementation of cloud services. For assuring the performance, availability, security and scalability of cloud services and its components, such as data center infrastructure, application performance and security, third party partners like Spirent Communications provide a wide array of testing solutions, including virtual test equipment not offered by any other vendors, along with qualified test team expertise, to deliver holistic dynamic data center testing solutions.

Source: Spirent

[8]Gartner RAS Core Research Note G00170153 Mix in the Right Test Skills to Achieve Quality, Thomas E. Murphy, 08 Oct 2009
[9]Gartner RAS Core Research Note G00174889 Survey Shows High Interest in RTI, Private Cloud, Donna Scott, 26 March 2010

About Spirent Communications

Spirent Communications is a leading technology company focused on delivering innovative systems and services to meet the testing needs of its diverse base of global customers. As a worldwide provider of performance analysis and service assurance solutions, Spirent enables the development and deployment of next-generation networking technologies such as broadband services, Internet telephony, wireless and Web applications and security testing. Working behind the scenes, Spirent ultimately helps the world communicate faster, better and more often. Market leaders offering a wide range of Internet and telecom-based products and services rely on Spirent's lab test solutions to evaluate the performance of emerging and existing technologies.

As new communication systems and applications are introduced, Spirent provides the tools for service management and field testing which improve the troubleshooting process as well as the quality of critical networks while enabling large companies and governments to secure and manage their networks. Spirent's proven record in helping the biggest equipment manufacturers, service providers and organizations with enterprises to succeed is built on the expertise and insight gained as a pioneer in testing virtually every major communication protocol, standard and technology.

Spirent also provides services expertise for all major communications vendors. Spirent's test methodology and automation experts include IETF engineers and RFC authors who have developed many of the industry's acclaimed benchmarks. They help define standards and performance tests for dozens of industry forums and standards bodies.