SOLITAIRE INTERGLOBAL

Comparing Virtualization Alternatives – What's best for your business?

A quantitative analysis of the business differentiators among x86, Unix, and System z virtualization technologies

1. Introduction

Virtualization on an enterprise level has developed into a significant strategy for organizations that are watching costs, but do not want to adversely impact service levels. The increasing need for agility in market response is also pushing more and more organizations to implement virtualization on an organizational level, with more and more production VM images being deployed every day. Virtualization provides both an isolation and prioritization of resources that allows a single platform to function as if it were split into multiple machines. The conjunction of today's technology-driven business marketplace with the economic clime pushes organizations into a continual search for higher efficiencies and better leveraging of IT resources.

Virtualization is one of the most powerful tools in the achievement of increased leverage and efficiency of those resources, while positioning organizations strategically for a cloud-computing model. The choice of virtualization method and platform can be challenging, as businesses struggle to understand the change in challenges to their information delivery processes, support staffing and the different, critical decision elements that need to be considered. Since the impact of virtualization forms an underlying contribution to an organization that is a diffuse layer within the IT infrastructure, IBM engaged Solitaire Interglobal Ltd. (SIL) to conduct surveys, gather data and perform analysis to provide a clear understanding of the benefits and relative costs that can be seen when organizations implement IBM z/VM as part of their IT architecture. This analysis has been primarily directed at the value of virtualization from a business perspective, so that those whose role it is to provide business leadership can understand the benefit of the IBM z/VM virtualization offerings when evaluating its selection.

During this study, the main behavioral characteristics of software and hardware were examined closely, within a large number of actual customer sites (79,360+). All of these customers include organizations that have deployed virtualization as part of their production environments. This group has organizations that maintain both single virtualization standard and those that allow a heterogeneous mixture of virtualization methods and mechanisms. The information from these customer reports, and the accompanying mass of real-world details is invaluable, since it provides a realistic, rather than theoretical, understanding of how the use of different types of virtualization can affect the customer.

In the collection and analysis of this data, a number of characteristics were derived. These characteristics affect the overt capacity, efficiency and reliability of the environment and its affects on operational and business performance. The behavior represented by these characteristics has then been projected and modeled into possible options for deployment. In order to build this understanding more than sheer performance is required. Although the performance of the virtualized systems is an important metric, the translation of that performance into business terms is more germane to today's market. The business perspective encompasses a myriad of factors, including reliability, security, staffing levels, time-to-market (agility) and other effects. This ties directly into the decisions that IT managers, CTOs, project managers and business leadership have to make daily.

2. Summary of Findings

The purpose of this analysis was to examine the real-world impact on businesses that deploy IBM's z/VM virtualization product, compared to those using UNIX or x86 products. For the purposes of this analysis, the different variants of x86 virtualization have been grouped together and treated as a single entity. Likewise, the different UNIX variants have been grouped. This encompasses PowerVM, Oracle VM for SPARC, and so on. If information on the relative differentiation among those groups and x86 is desired, a previous SIL study should be referenced. It can be found here.

The metrics used to analyze the differences in platforms were both objective and subjective. The objective metrics include reported data points on costs, run times, resource usages, and so on. The subjective metrics include responses on various levels and sources of customer satisfaction and perception. While overall customer satisfaction uses a variety of qualitative and quantitative measures, it still provides an end-result measurement of deployment success for the customer. A few of the highlighted findings can be seen in the quick summary below.

Quick Summary

Category	Commentary	Quick Byte
Customer Satisfaction	The more complex or volatile the environment, the more all aspects of customers reported high satisfaction with z/VM.	z/VM provides enterprise-level support for complex and changing customer needs.
Total Cost of Ownership (TCO)	While TCA can be more for z/VM than some of the alternatives, the TCO rapidly changes that picture, especially when a multi-year view is taken.	TCO can be as little as 17.69% of the expense of competitive offerings.
Staffing	The normalized staffing levels for z/VM are smaller than those for the competitive offerings when a complex environment is needed by as much as 13 *times* less than other options.	Staffing is far more effective with z/VM – the same number of people can support up to 3 *times* more VM images.
Risk	The reported risk of deployment is considerably better for z/VM users, with competition incurring as much as 11 *times* more risk to operations.	Enterprise resiliency from z/VM safeguards the IT operations.

Category	Commentary	Quick Byte
Availability	The more virtualized the environment, the more critical the availability becomes. z/VM requires fewer platform and VM reboots than competitive platforms with very little downtime, either planned or unplanned.	Enterprise strength availability from z/VM supports global operations and midnight sales.
Agility	z/VM users are reporting faster deployment times by as much as 5 times, and with much more predictable delivery of the new systems and features.	A well-managed z/VM system is directly associated with faster and predictable time-to-market.
System Efficiency	With the ability to push the utilization of the z/VM resources to substantially higher levels than any other platform, the z/VM platform routinely supports more than 17 times more VM server images than the competition.	The z/VM platform allows more IT workload to be done with fewer overall resources.
Security	z/VM supports all forms of security control and isolation, including those required for highly secure implementations, separating resources for memory, network, I/O and access.	No reported successful VM hacking in z/VM means organization assets are well protected.
IT Service Quality	The ability of IT to support the business with reliable, safe and consistent services forms the base for quality. z/VM service quality is as much as 5 *times* higher than the other options.	z/VM delivers enterprise-level quality IT services for all organizational business needs.

These key findings are all substantial reasons to consider z/VM for an organization's virtualization and architectural choice. In addition to the key findings, an analysis of the relevancy of the mainframe architecture in today's business-case driven market provides a critical view of how it has changed in reference to the System z and z/VM architecture, showing that most of today's businesses have a true business case for mainframe deployment.

2.1. Study Scope

In order to understand the impact of IBM z/VM and virtualization as a key part of an organization's IT, a large number of deployments were examined. These deployments included situations where the virtualization choices were homogeneous within an organization and ones where a mixture of different methods, software tools and components existed. The relative degree of difference in operating behavior for each factor, i.e., total number of outages, etc., was then compared to understand the net result of the respective combinations. The effects were observed in general performance and capacity consumption, as well as other business metrics.

2.2. Methodology

The approach taken by SIL uses a compilation and correlation of operational production behavior, using real systems and real business activities. For the purposes of this investigation, over 79,360 environments were observed, recorded and analyzed to substantiate the findings. Using a large mass of customer and industry experiential data, a more accurate understanding of real-world behavior can be achieved. The data from these systems was used to construct a meaningful perspective on current operational challenges and benefits. The reported behavior of the systems was analyzed to isolate characteristics of the architecture from both a raw performance and a net business effect perspective. This information was then projected on the production system performance of the non-z/VM deployments to better understand the possible impact and effects. All input was restricted to those organizations using operating systems in versions that were current in calendar years 2011 and 2012. Since many of the components in this environment have releases at staggered points in time, only those components that were either the current version or a -1 version based on those calendar restrictions were included in the study. Additional information on the methodology and study diversity can be found in the supplemental methodology notes at the end of this document.

In a situation such as that presented by this study, SIL uses an approach that incorporates the acquisition of operational data, including system activity information at a very detailed level. It should be noted that customers, running on their production platforms, provided all of the information. It is essential to understand that none of the data was captured from artificial benchmarks or constructed tests, since the value in this study comes from the understanding of the actual operational process within an organization, rather than the current perception of what is being done. Therefore, these sites have tuning that is representative of real-life situations, rather than an artificial benchmark configuration. Since the focus of this analysis was not to tightly define the differences among the different minor variations of operating system or hardware, the various releases were combined to show overall architectural differences. This provides a more general view of architectural strategy.

The study was further restricted to organizations that have larger implementations. While this restriction is not intended to make a statement on suitability of any virtualization mechanism for small organizations, it is true that smaller processing demands are more easily handled, and provide smaller differentiation in analysis. For these reasons, the smaller implementations[1] were filtered from the study.

The information in this study has been gathered as part of the ongoing data collection and system support in which SIL has been involved since 1978. Customer personnel executed all tests at SIL customer sites. The results of the tests were posted to SIL via the normal, secured data collection points that have been used by those customers since their SIL support relationship was initiated. As information was received at the secure data point, the standard SIL AI processing prepared the data in a standard format, removing all detailed customer references. This scrubbed data was then input to the analysis and findings.

The analysis of this data has produced findings in two groups of viewpoints – business management and technical. For a more concise summation, those findings have been discussed separately in the body of the paper.

[1] *The guidelines for organizational size classification that SIL uses are defined in the supplemental methodology notes at the end of this document.*

2.3. Business Perspective

Ultimately, IT and technology are designed to support business functions. So one of the primary perspectives of the study was the view of the technology by an organization's business management, both executive and line-of-business. For the purposes of this part of the analysis, the patterns of operations from the study organizations have been grouped into similar categories and then compared to identify their influence on business metrics.

These metrics are:

- Customer satisfaction
- Total cost of ownership
- Staffing
- IT stability and reliability
- Agility (time-to-market)

Each of these business metrics has measurable and significant differentiation when the projected IBM z/VM deployment solution is viewed.

The more granular business metrics are those measurements that show how a specific success criterion is different in the general population of the implementers versus those that have deployed z/VM. For further clarification, those situations where UNIX- and x86-based virtualization mechanisms were deployed have also been broken out. These metrics are fairly broad in coverage and touch on areas of financial consideration, as well as organizational quality. They are presented with short definitions and the focused net effect of IBM z/VM deployment. In order to be meaningful across a variety of industries, all of them have been normalized on a work-unit basis[2], and categorized by levels of organization size (medium, large and very large). The base measure has been set by the medium company average, so that all other metrics are based on a variance from that standard set point. The implementations included in this study have been restricted to those implementations in production.

Customer Satisfaction – End IT User

The ultimate measurement of a successful implementation is *customer satisfaction*. SIL tracks this metric split out among the end IT user, IT operational or line-of-business management, and executive management from each organization, since the perception of satisfaction may radically differ among those groups. The satisfaction of the end IT user about their computer systems tends to focus on the delivery of services, rather than the virtualization itself, although no application can successfully work effectively with a poorly configured or fragile virtualization method. That being said, the satisfaction with IT implementation and operation provides the most general metric for evaluation. This satisfaction rating was obtained from a large group of customers and provides a singular perspective on the overall success of virtualization deployment. While this is a subjective rating provided, it does provide the business' actual perception of success.

[2] *Work-unit basis has been defined using the published International Function Point User Group standards and are based on function point (FP) analysis.*

The advantages seen by the reporting clients show increasing satisfaction in the applications run under z/VM, much of which can be attributed to the number of tracked complaints from end IT users of those systems. The following chart shows the reported average monthly complaint count for the different platform groups. These complaints have been restricted to continued operational issues, and exclude complaints associated with missing and desired application functionality.

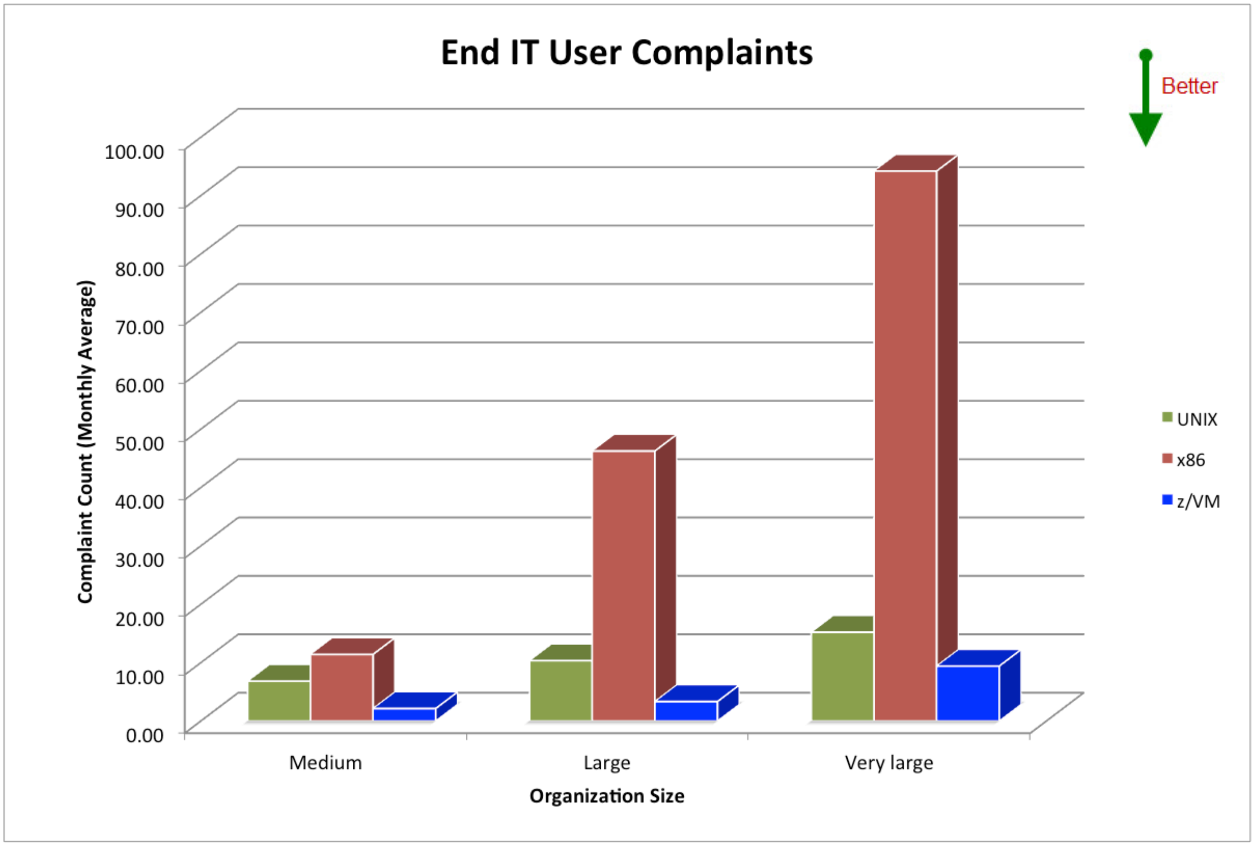

The three top reasons cited by reporting customers for the satisfaction were:

1. Smooth running operation with little downtime
2. Consistent and dependable run times
3. Speed of implementation

Over 99% of the respondents in this study cited one of the three reasons listed above.

Customer Satisfaction - Operational

The operational perception of the customer, based on a variety of component metrics (e.g. support levels, communication, price, etc.), demonstrates satisfaction and success at the most generic level. This satisfaction metric is different from the overall satisfaction metric described earlier, in that the previous metric was gathered from the executive management level, while this metric examines the feedback from the operational side of the organization. This specific metric comes from information reported both by the IT departments and the line-of-business (LOB) groups.

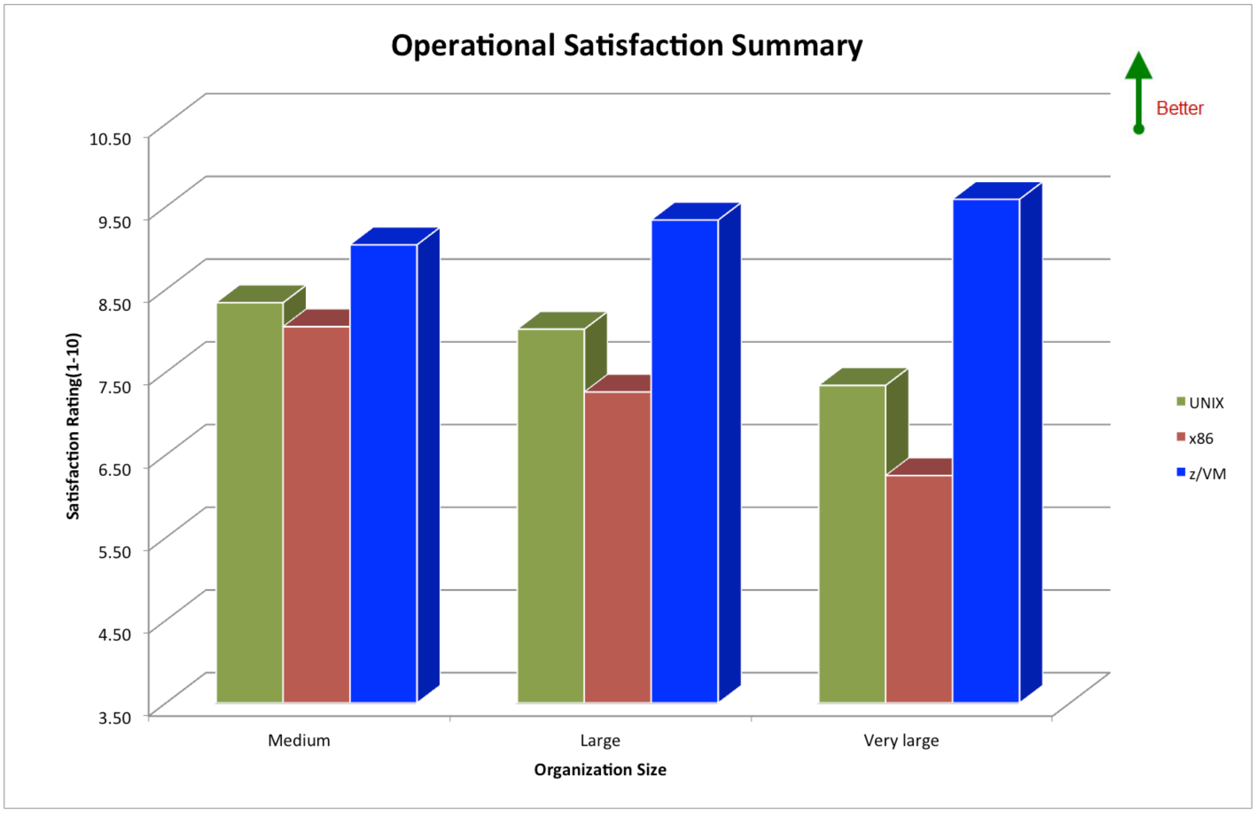

The satisfaction of the IT operational staff and the LOB with the z/VM deployments reflect the reliability and resiliency of the platform as a deployment choice, in addition to the previously mentioned integration benefits. The most highly cited reasons for the satisfaction were:

1. Smooth running operation with little downtime and complaints
2. Efficiency of operational support
3. Large amount of information available to manage operations

More than 96% of the reporting customers cited one or more of these three reasons for their satisfaction.

Customer Satisfaction – Executive Management

The final place for verification of system success is in the executive management position. The satisfaction of the customer executive management with their IT systems tends to focus on the application and cost, rather than the virtualization. An overall summary of the satisfaction of the executive management of the studied organizations shows some interesting patterns.

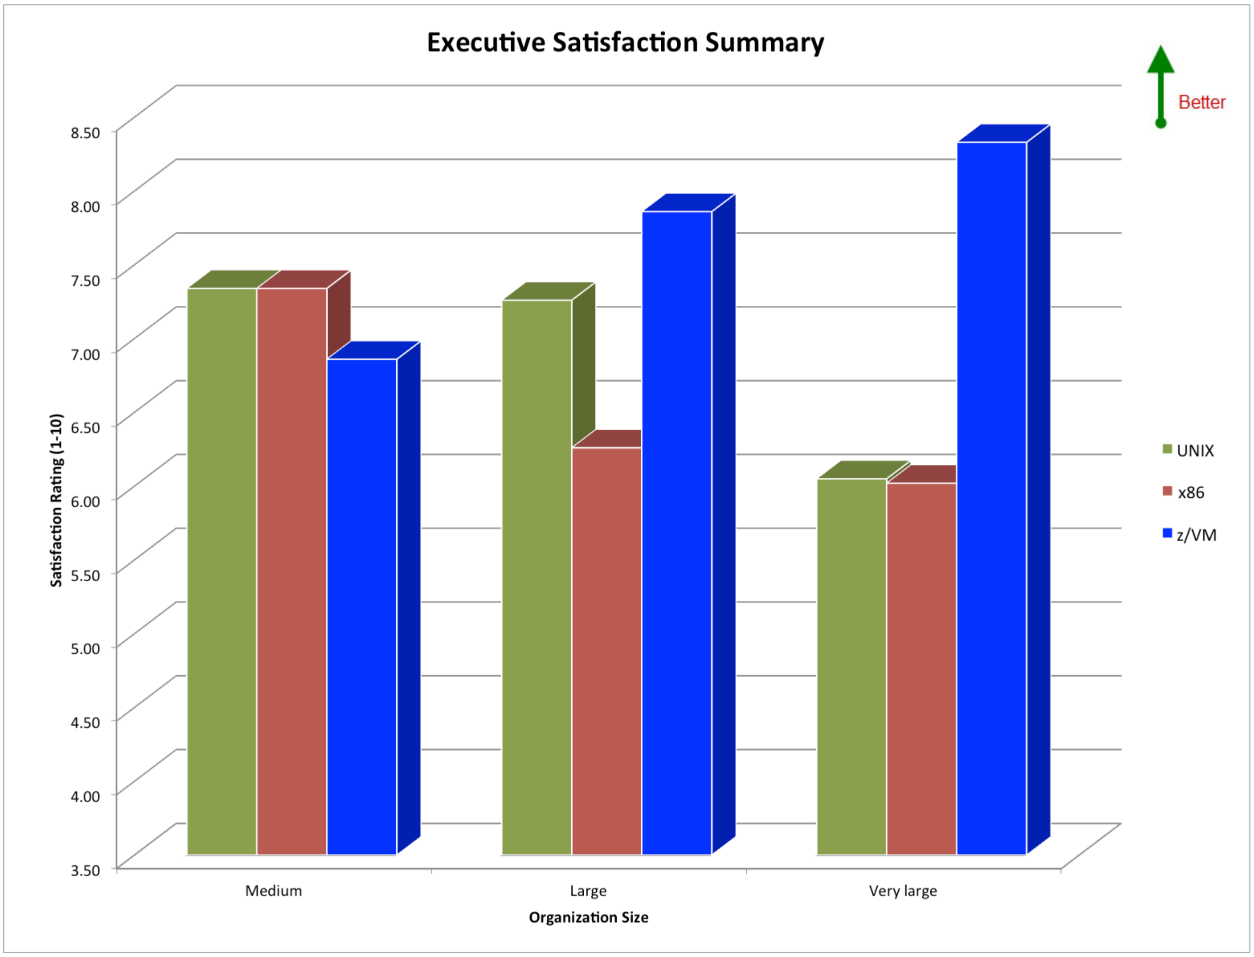

Of course, one major source of the perception of success by organizational executives is the number of complaints that they receive about the system operation. As seen in an earlier section, z/VM displays a significantly lower number of complaints. While the specific customer complaints can be affected by management techniques, application design and other factors, the relative comparison is a legitimate indicator of how well the operating system supports the processing at the organization. The three top reasons cited by reporting customers for the satisfaction were:

1. Smooth running operation with little downtime and complaints
2. Low staffing
3. Speed of implementation on System z

Since one of the main reasons for executive satisfaction is financial, further analysis was done to examine the degree of utilization, or efficiency, of each platform. When this was included in the analysis, it is obvious that in those cases that platforms are highly leveraged, the z/VM satisfaction increases substantially.

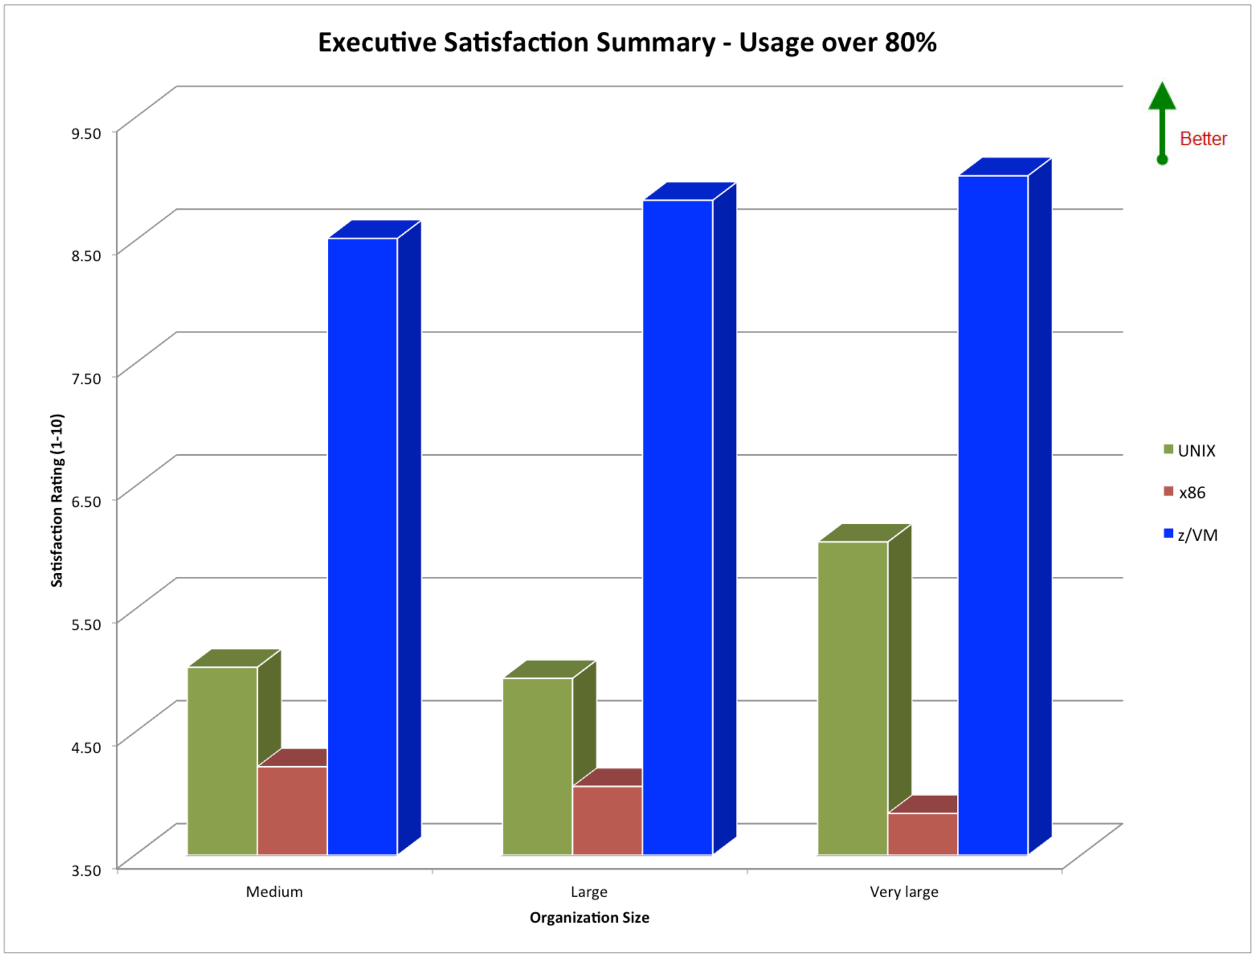

It is very understandable that a low utilization of a substantial resource, such as a mainframe computer, would be a source of dissatisfaction with executives. However, the high degree of satisfaction for the efficient load and management of that resource shows that it is more of an organization strategy rather than a platform limitation.

> **Research Note:** The study shows a very high correlation between the system efficiency and the satisfaction of executive management. It also suggests that an evaluation of organizational business process and strategic resource management would be a reasonable approach for increasing overall satisfaction.

Overall Expense (TCO)

This cost perspective looks at the total cost to the corporation during a specific time period. This is normalized on three bases: employee, sales revenue and legal entity count, and contains expenses associated with up to a 3-month deployment preparation phase. These expenses span all of those included in the operational cost metric and are supplemented by expense contributions for physical plant, corporate overhead, long-term investments, etc. The TCO financial metric is more comprehensive than a straight operational metric. This metric it should not be viewed in isolation, since extraordinary expense patterns for individual organizations may cause minor variance in the exact comparison values. For this reason, the comparison metric should be viewed as indicative and providing a general range rather that an exact value. However, with the large number of contributing organizations, the data is sufficiently large that, combined with the other business metrics, this comparison helps to set an appropriate perspective.

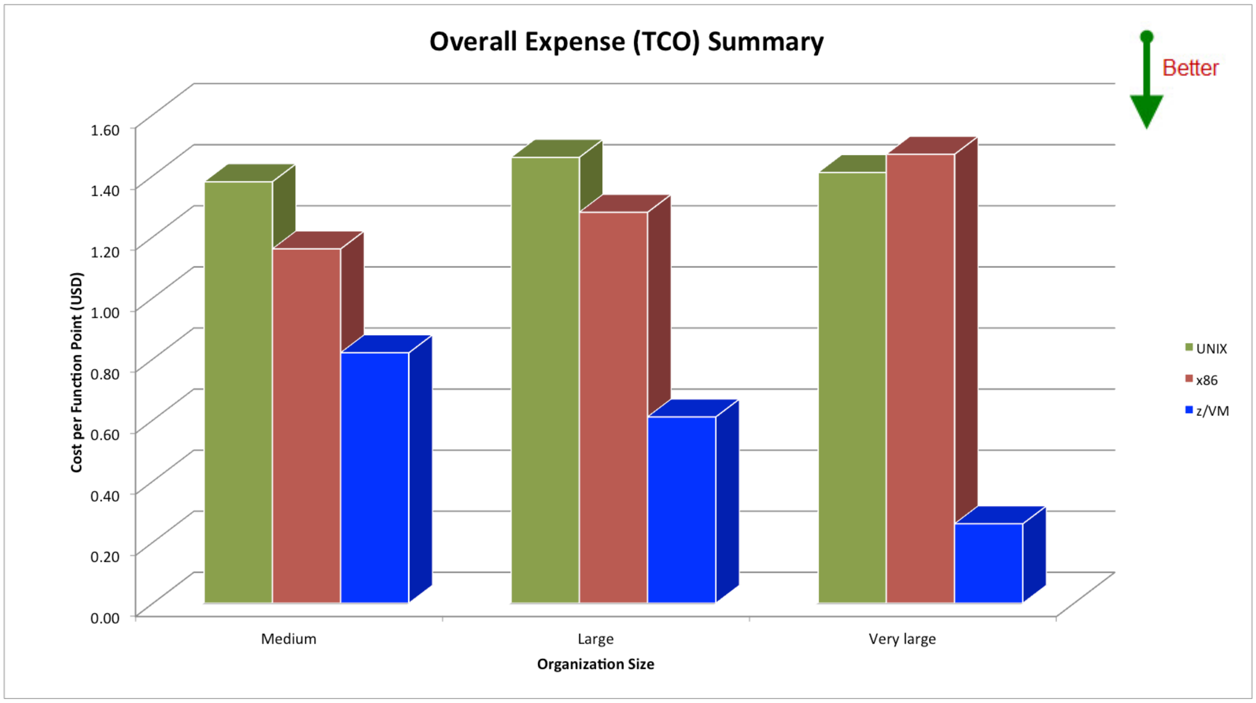

The IBM z/VM application show lower overall expenses (as much as 82.31%) over a wide range of organization size, although the smaller organizations do not see the efficiencies of scale that the larger organizations see with the z/VM product. It should be noted that this TCO comparison should be viewed in conjunction with availability and downtime metrics. Since no cost has been associated with unavailability, each organization should factor in its associated downtime cost to the TCO metric provided here. The downtime metric can be found later in this document.

Additionally, the cost of acquisition is higher with the System z platforms than for the smaller Intel and UNIX platforms. This disparity in cost levels is obviated when the level of virtualization and capacity demand increase. This switch in the defining metric from TCA to TCO happens in all situations eventually, but is more rapid in the larger deployments. Since the TCO holds true as a metric, well past the usefulness of the TCA, the TCO has been used as the defining cost metric. The differential among the solutions is based largely on the lower expenses for the efficient deployment and the lower overall cost of the solution, including staffing. This is affected strongly by the scope of the virtualization deployment, with increased expenditure efficiency present as the complexity and size of the virtualization deployment increases. Customers of all degrees of deployment reported a consistent pattern of differentiation in three main areas:

1. More highly-leveraged platforms
2. Lower datacenter costs (environmental, facility, etc.)
3. Lower staffing costs overall (due to tools, stability, etc.)

An interesting metric can be seen if the efficiency of scale (EOS) is examined for the virtualization options in this area. This measurement looks at the change in the normalized cost as the implementation increases in size and complexity in either the physical deployment or the number of VMs. It reflects any efficiency that tools and management flow provide in a specific virtualization mechanism. Using the organization size as the driving principal, the TCO EOS trend can be summarized as shown in the graph below:

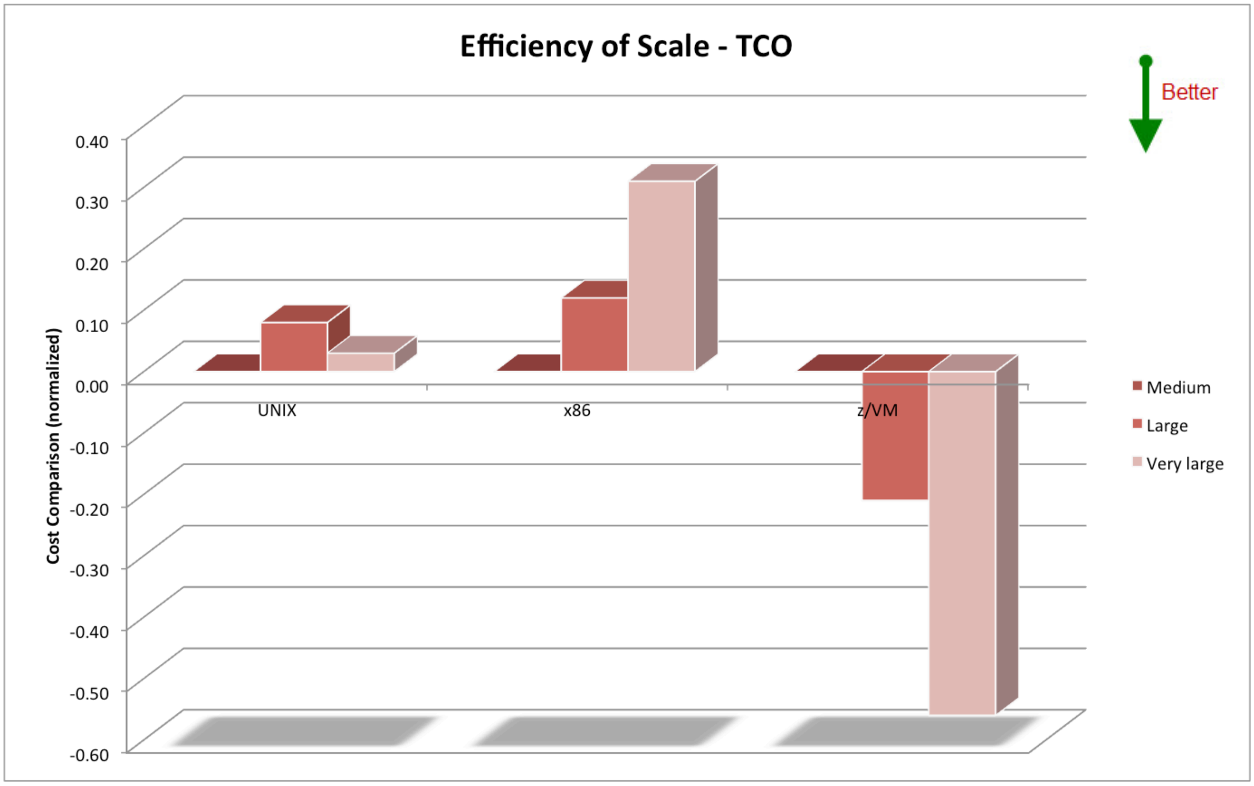

It is notable that the z/VM data shows a clear implementation of the efficiencies of scale, which is counter to the competitive offerings. The top three sources of this advantage were reported as:

1. Speed in production deployment
2. Efficient resource sharing
3. Stability of platform integration (storage, etc.)

It should be noted that there are substantial differences in the specific x86 and UNIX variants in this area. If either of those two architectures is of interest, more detailed information on the specific VM products can be seen in the paper referenced at the beginning of this study, which focuses in specifically on virtualization offerings in those two architectures.

These factors have produced a realized savings in the cost per VM of about 68% for z/VM, when the complexity and size of the virtualized environment moves from medium to very large, while the competitive offering actually grow in cost per VM, up to 1.35 *times*. This efficiency can be substantial when viewed for the enterprise, saving millions of dollars in deployment and operational costs.

Staffing

An underlying factor that shows itself in many other areas is the effectiveness of the interface between the technical user and the infrastructure, including software, hardware and operating system components, and the subsequent effect on staffing. The efficiency of any of the specific components that provide that influence on the user experience are difficult to break down into metrics other than in overly-detailed comparisons that lose their effectiveness by virtue of the degree of detail. Therefore, a general view of the full-time staff position equivalents was reviewed to provide a general metric for the platform comparison. These levels are those required to maintain a "gold standard" environment for each operating system group. Once again, in order to provide a level comparison field, the workload on the systems was normalized to identical levels. The set point for comparison was selected as the staff level for a medium-sized organization using VMware.

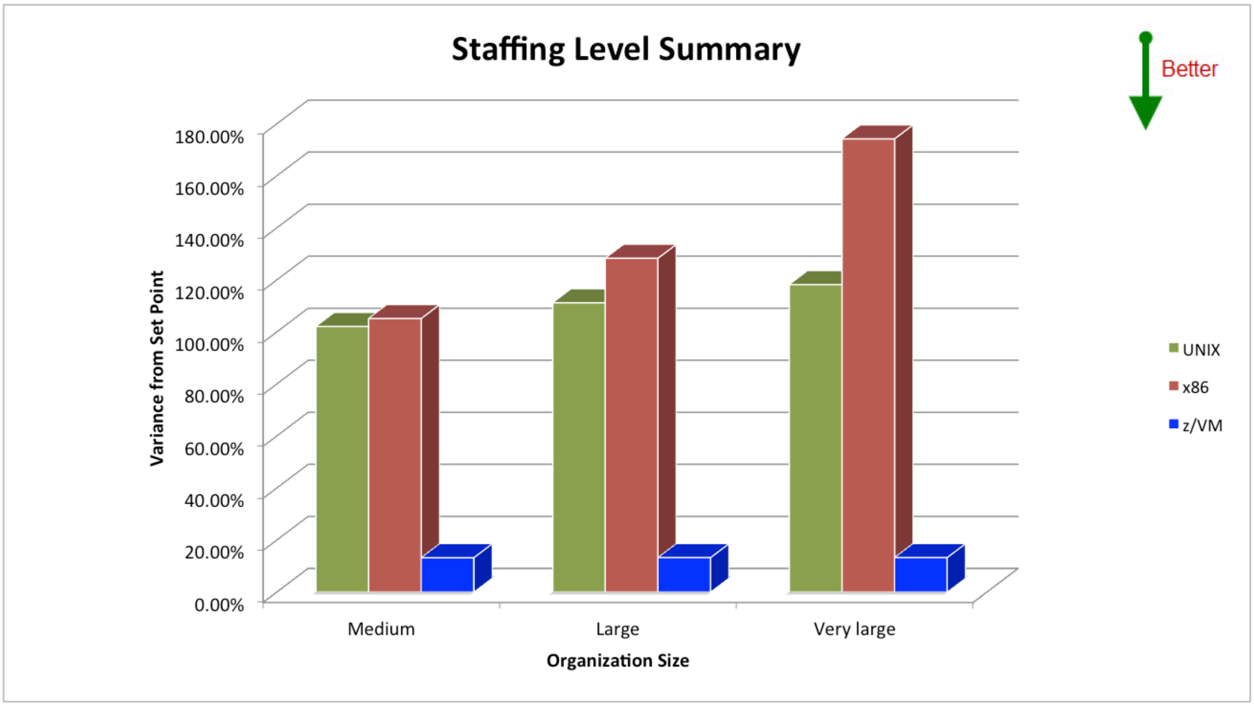

Since different virtualization methodologies have varying sets of implementation standards, it is important to keep the rigor of those standards in mind when reviewing the staffing. The noticeably lower staffing level for z/VM deployment and use is directly attributable to an efficient unified workflow, as well as a substantially different and fully integrated mechanism to handle the allocation of virtualized resources. This is of special note as the organization increases in size or if an organization is on the path to a cloud service delivery model. The normalized staffing levels for z/VM are smaller than those for the competitive offerings by as much as 13 *times*.

Another way to examine the staffing requirements is to look at which areas of activity consume the staff hours. For the purposes of this analysis a subset of the reporting organizations allowed SIL full time-motion data. This data was then analyzed to build a list of the top activities that the staff supporting virtualization performed. This occurrence analysis uses the frequency of the action to determine the weighting.

Task Frequency Summary

Rank	Description
1	Check resources levels
2	Reallocate and prioritize
3	Setup VM
4	Tune performance
5	Move applications - across environment types
6	Setup new server and VM
7	Move applications - within environment type
8	Install patches and fixes

From a time perspective, the task list order changes, since some of the frequently performed tasks are simple and quick, while some of the other tasks take a considerably longer time. The ranking in this table is in order with the most time-consumptive task first, the second next, and so on.

Task Duration Summary

Rank	Description
1	Move applications - across environment types
2	Tune performance
3	Move applications - within environment type
4	Setup VM
5	Reallocate and prioritize
6	Setup new server and VM
7	Install patches and fixes
8	Check resources levels

If the top three task areas are examined from a relative time consumption perspective among the virtualization options, an interesting pattern appears. This summary compares a normalized environment against the set point, which in this case is the one set by VMware at a medium-sized organization.

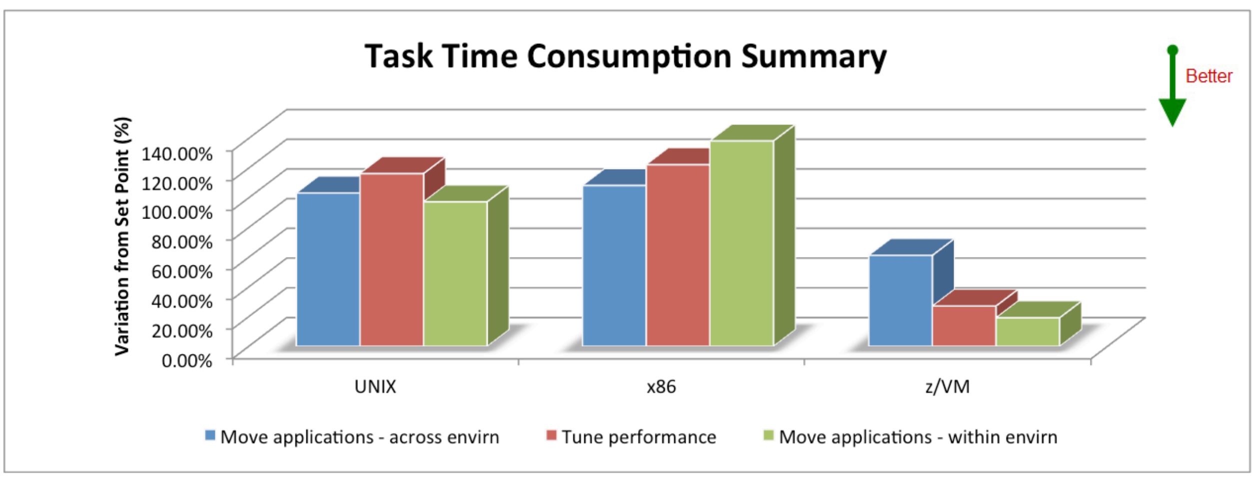

There is a radical difference in the amount of time spent on the top three staff time usage tasks when z/VM is included in the analysis. The z/VM advantage is as much as 7.19 *times* in these most heavily performed staff tasks. Part of this difference can be correlated to the workflow design within the z/VM management tools. The overall context switching was significantly smaller (94.2% less frequent) than the average. What this means is that an IT support person performing virtualization tasks has to switch workflow direction, opening additional screens or recording information to then change the open action on their screen 95% less frequently. This makes for fewer mistakes and faster task completion. This is supplemented by extensive integration within the toolset, operating system and hardware. This architectural and philosophical difference has the effect of further reducing the actions required from the supporting personnel by another 26% on average. All of this means that z/VM requires much less personnel time to accomplish the same operational result within the virtualized environment.

This can also be seen in the reported learning curve timeframes. The data from the reporting organizations included the interval of time that a staff member needed to be fully functional in each of the virtualization products. This was not a timeframe for expertise, which has many different metrics, but the base one of adequate performance, obviating the need for training supervision. This information is shown in graphic form below.

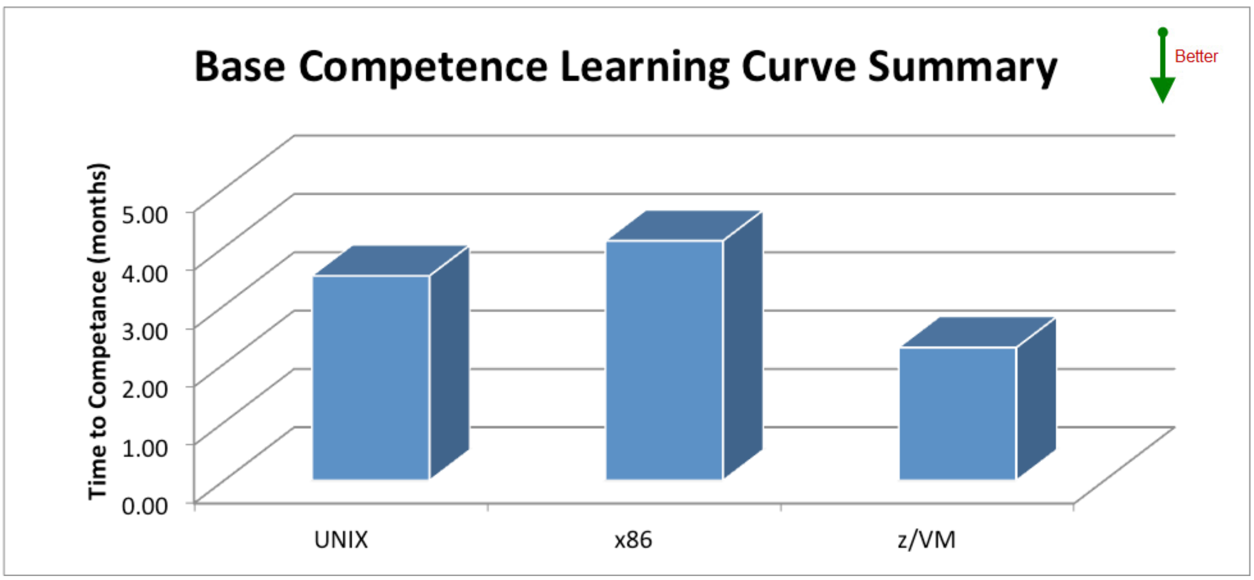

The learning curve on z/VM takes a significantly shorter time to competence than UNIX or x86 virtualization methods, with the most frequently stated reasons of:

1. Robust operational and management tools
2. Wealth of monitoring data and meta-data available
3. Optimized workflow

The faster ramp-up time of the z/VM virtualization method is as much as 1.8 *times* faster than the others in this study. This faster time to competence can be critical for organizations, as they deploy new virtualization efforts.

An interesting side note to this point in the analysis is that there was a substantially lower number of new z/VM staff FTE that reported into this study than for the competitive offerings. This was not the result of significantly fewer virtual partitions, but instead was because there is such a high degree of efficiency with the z/VM platform, that the existing staff personnel could handle the workload without supplemental resources. To better understand the thresholds of support, data from all of the organizations within this study were summarized. The thresholds of the number of VMs that a single FTE could support are shown in the chart below, with some significant splits. The reason for that granularity is that during the study process there were many observations made on overall system load, and the associated thresholds and efficiencies of the various architectures. One of the areas that formed a consistent thread throughout the metrics, including satisfaction, costs and others was the threshold of support personnel. This threshold looks at the number of VM images that are supported in actual operation. Although the normal SIL view of operations filters activities to only production images, in this one area the effect of the non-production environments is significant. Therefore, the chart below shows both the production-only view of support as well as the total support environment, including production and non-production.

The chart also includes two slightly different perspectives on the thresholds. The first is the total average of VM images across all reporting organizations. This tends to smooth out variations among organizations, staff members, etc. While this is a useful metric, in any organization there are individuals that are concentrated on VM support and others that perhaps are less leveraged. The category labeled “maximum” is the average of the most heavily loaded person at each organization. This provides a high-water mark threshold average that is in some ways more informative than the overall average number, and factors in differences between dedicated personnel and those that are more part-time.

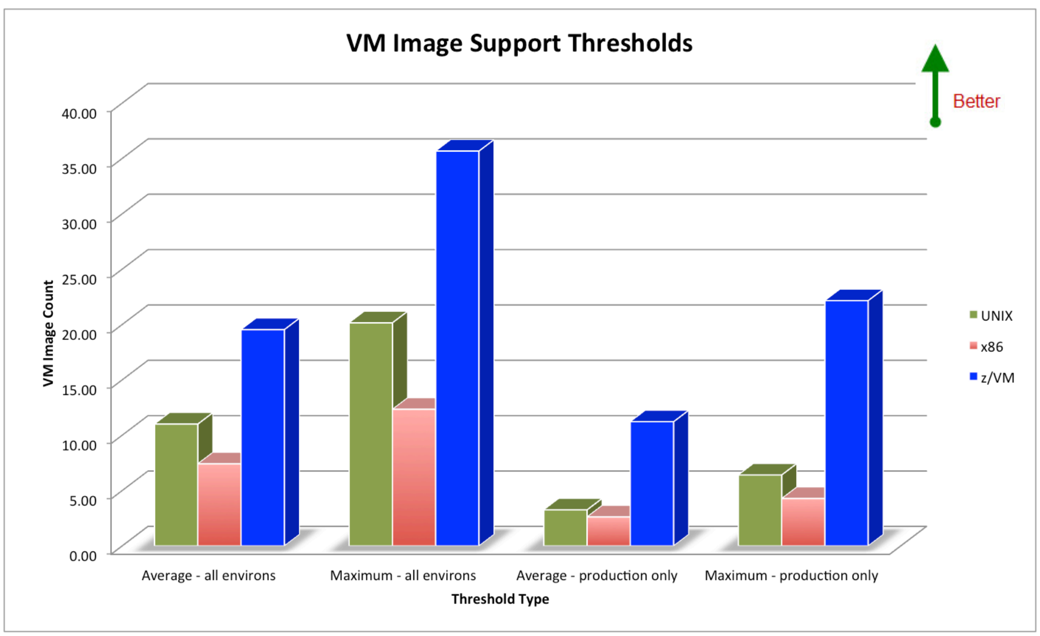

The influence of the platform integration with the VM product results in significant differences in the support thresholds. The z/VM option has a demonstrated threshold that is more than 3 *times* higher than the competitive offerings for total production and non-production environments and more than 5 *times* higher than production VM images alone.

IT Stability, Risk and Reliability

Risk is composed of many factors. It includes the stability and reliability of the platform, as well as the chances of platform failure. IT stability and reliability metrics include all downtime, both planned and unplanned. The dependability of the implementation is a combination of the individual reliability of each component, along with the quality and effectiveness of the actual implementation. As such, both the planned and unplanned outages affect the overall usability of the total system. SIL views availability as a combination of all outages, i.e., network, hardware, OS, DBMS, etc. The number of outages has been normalized for a 20-virtual production partition operation, with both planned and unplanned outages included. Where virtualization has been included in the architecture, each of the virtualized environments has been considered as a separate platform. Each of these outages takes valuable access time away from the corporate resources. The following chart shows the percentage of time that those outages represent and includes all forms of unavailability, irrespective of source.

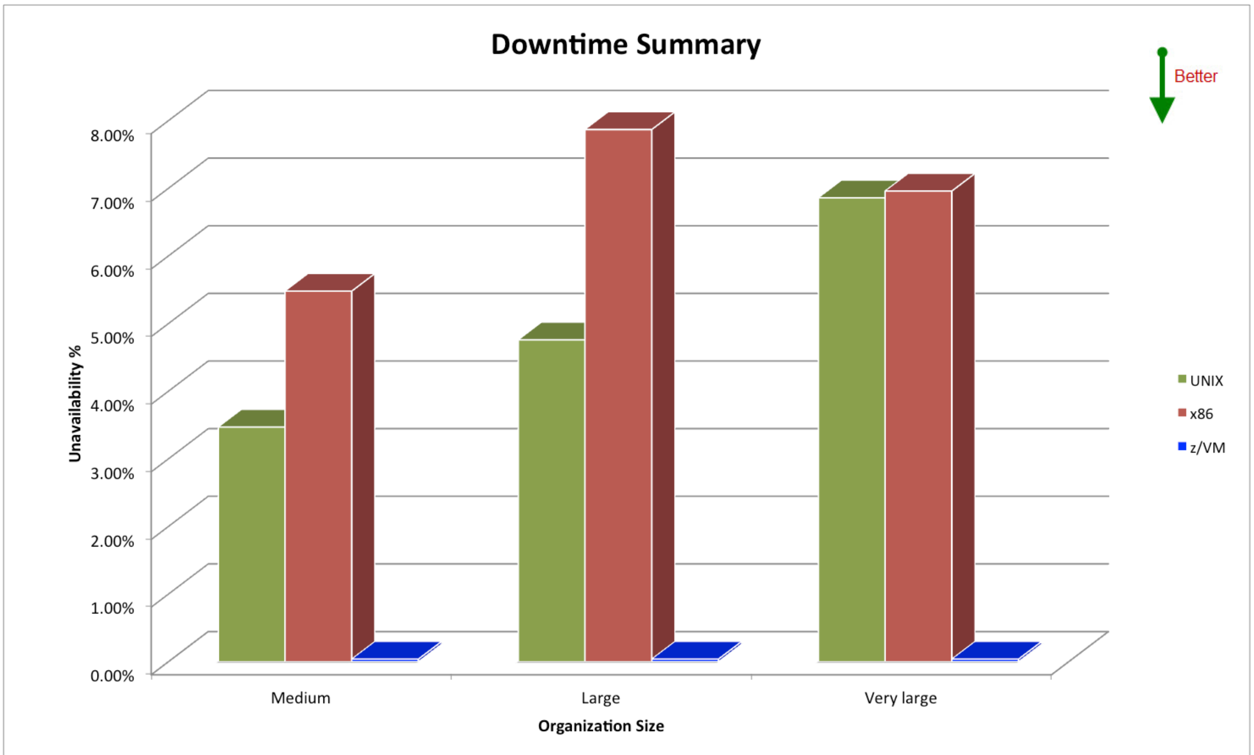

As shown above, there is a substantial indication of how the z/VM virtualization contributes to both stability and reliability of an organization's implementation, due to the combination of high performance and native resilience. The three most cited sources of the high availability from customers are:

1. Ability to automatically move resources to needed processes
2. Limited need to reboot the full platform
3. Fewer system patches and updates required

It should be noted that since SIL uses actual production data, recent significant losses of system uptime on highly publicized virtual environments have contributed to the lower availability numbers. There are also significant variations in uptime and availability among the different specific UNIX and x86 offerings. These are documented in the previous SIL study mentioned at the beginning of this document.

> **Research Note:** Recent data shows that many of the availability numbers for both x86 and UNIX are dropping for production virtual environments, as the complexity and workload of those implementations increase. The decay of availability is NOT present for z/VM implementations.

It should be noted that the practices of the individual organization when viewed from a best practices perspective makes a difference in the amount of planned downtime. However, the overall trend in availability is a definite indicator of platform stability.

The cost of that availability is difficult to articulate, primarily because such a cost estimate has significant subjective components. However, a quick analysis of the customer-reported financial impact of outages yields a general metric that provides some interesting insights.

SIL considers risk to be comprised of three components:

- Percentage chance of component failure
- Percentage chance of budget or timeframe overrun
- Potential exposure, expressed as a percentage amount of overall budget or timeframe overrun

These three percentages are added to form the overall risk factor for a scenario. The risk factor summary for the platform scenarios is shown below.

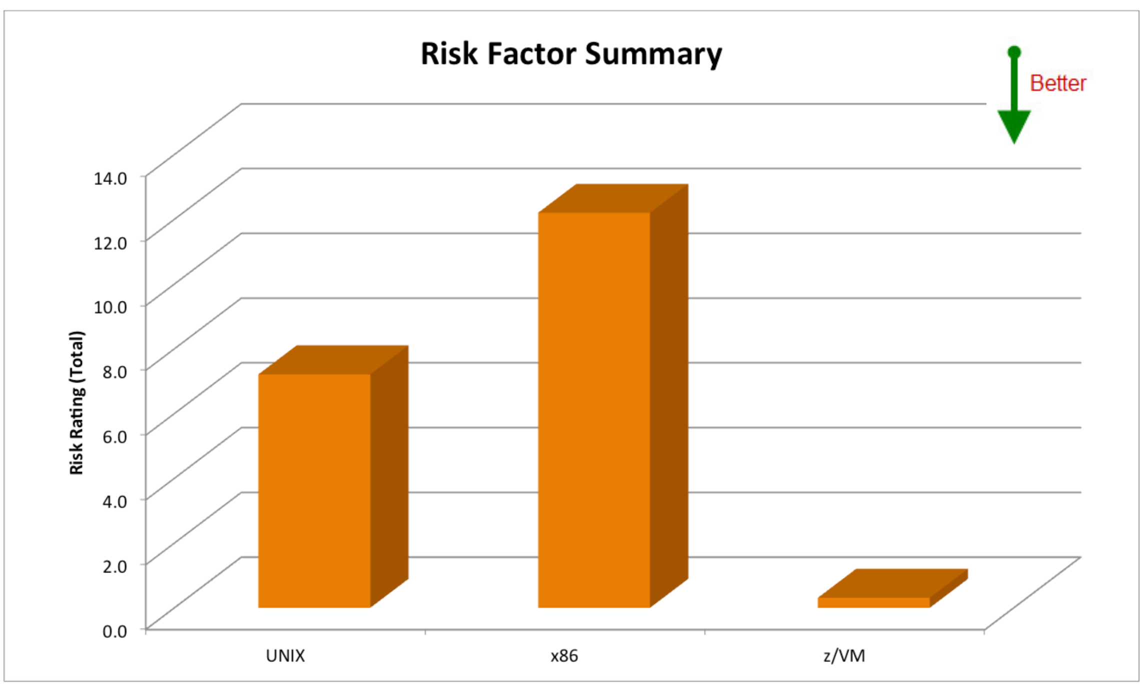

This graph shows that there is demonstrated risk mitigation from the general operations experience when using z/VM. The risk exposure for z/VM is significantly smaller than the competition, with z/VM deployments showing only 9% of the risk that has been reported for other virtualization methods. Much of this lower risk can be attributed to the high resiliency of the deployment and increased efficiency of the resource allocation within the virtualization component itself, which significantly lowers the risk of component failure.

Agility

Agility is defined as the average number of calendar days from the start of an initiative to the start of full production operations for a project. This is NOT staff days or hours, but the actual calendar span, including all weekends, holidays, etc. All of the contributory factors, such as staffing and reliability, radically affect the speed in which a company can move a business concept from inception to market. This nimbleness is a key element of increasing market share and continued corporate viability. While the performance metrics were gathered on the production systems, additional measurements were also collected to track the amount of time that the systems took to move from initial conception to full production implementation. The results demonstrate a significant increase in agility when platforms running z/VM-virtualized environments were used. This increase in agility has been reported to be as much as 5.1 *times* faster for the z/VM systems when compared against the overall study group. This translates into a faster time-to-market for business initiatives. The comparison is intended to be evocative and not quantitative, since other critical success factors, such as management methodology, resource availability, etc., can enter into this picture.

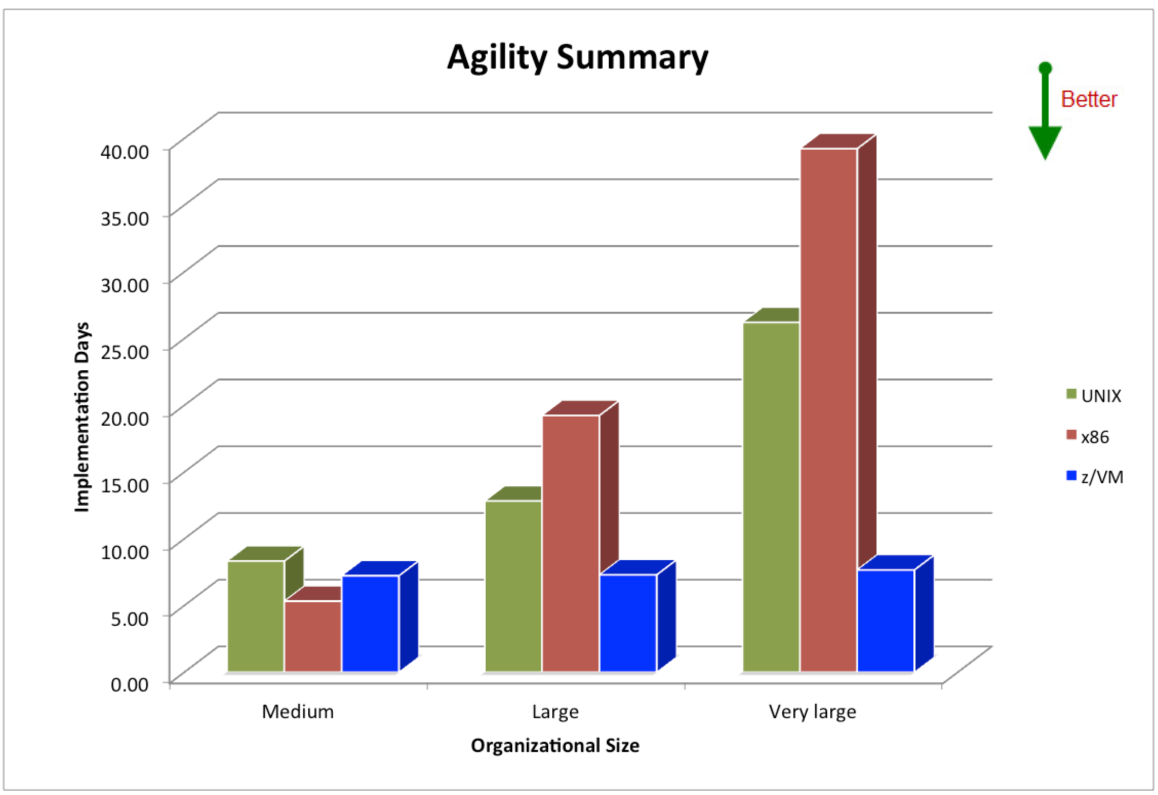

It is apparent from the reported data that there is a definite agility advantage to using z/VM-deployed systems as compared to the overall experience, especially when organization standards for production system promotion are comprehensive. When asked for specific sources of the agility, the most frequently cited reasons from customers were:

1. Speed of movement from non-production to production environments
2. Ability to easily shift resources to accommodate new implementations
3. Management tools and reporting mechanisms

The differences in agility can be substantial, with z/VM showing faster deployment times by as much as 511.2%. The faster time-to-market advantage that is present with the z/VM system can be a major competitive advantage for any organization.

2.4. Technical Perspective

One of the main perspectives for this analysis is from the viewpoint of the IT professional. Since IT needs to understand the underlying architecture and important characteristics of any technology, this perspective tends to focus primarily on the objective understanding of what a z/VM deployment can contribute and will require. This understanding encompasses some basic performance characteristics and operational challenges.

The more technical perspective still has pertinence for the business evaluation since it concerns those elements of operation that result in significant risk, performance and efficiency factors. The metrics that fall into this area are:

- Consistent delivery of data and application service
- System efficiency
- Environmental impact
- Application efficiency
- Management visibility and control
- Security
- Strategic Positioning for Cloud Deployment

These metrics form a picture of the operational side that is the more granular underpinning for the business support of virtualization.

Consistent Delivery

One of the original perspectives for this study was a view into the reliability of the various virtualization mechanisms. One of the metrics within that view is reliability. However, reliability can be somewhat limited in system views, since it is frequently only seen as uptime. This ignores situations where the platform is up, but the application is not available, as well as a whole other dimension of repeatable, consist delivery. Perhaps it is better to look at this area as quality – encompassing uptime, consistency and dependability. While this is not a common view of the contributions of a platform to IT operations, it is extremely important. If a platform is not up, available and providing consistent performance, it is very difficult to deliver quality services to the users within an organization. When the quality factors were examined for the reporting organizations, it was notable that the variance in runtimes among the platforms for ongoing equivalent business services and transactions varied widely by platform type. The z/VM consistency is high, with repeated execution variance running 3.1-6.2. The UNIX platform VM images had much wider swings in timing delivery – varying by 173.1%. The x86 platform VM images varied even more, with swings in execution time by as much as 416.2%. This means that a normal 37-minute business transaction process run on z/VM may take as much as 39.3 minutes, while the same process on a UNIX platform can take 101 minutes and on the x86 platform it can take 191 minutes. This consistency of delivery is important on many dimensions, since it sets not only customer expectations, but also the base for operational schedules, service level agreements and resource allocation.

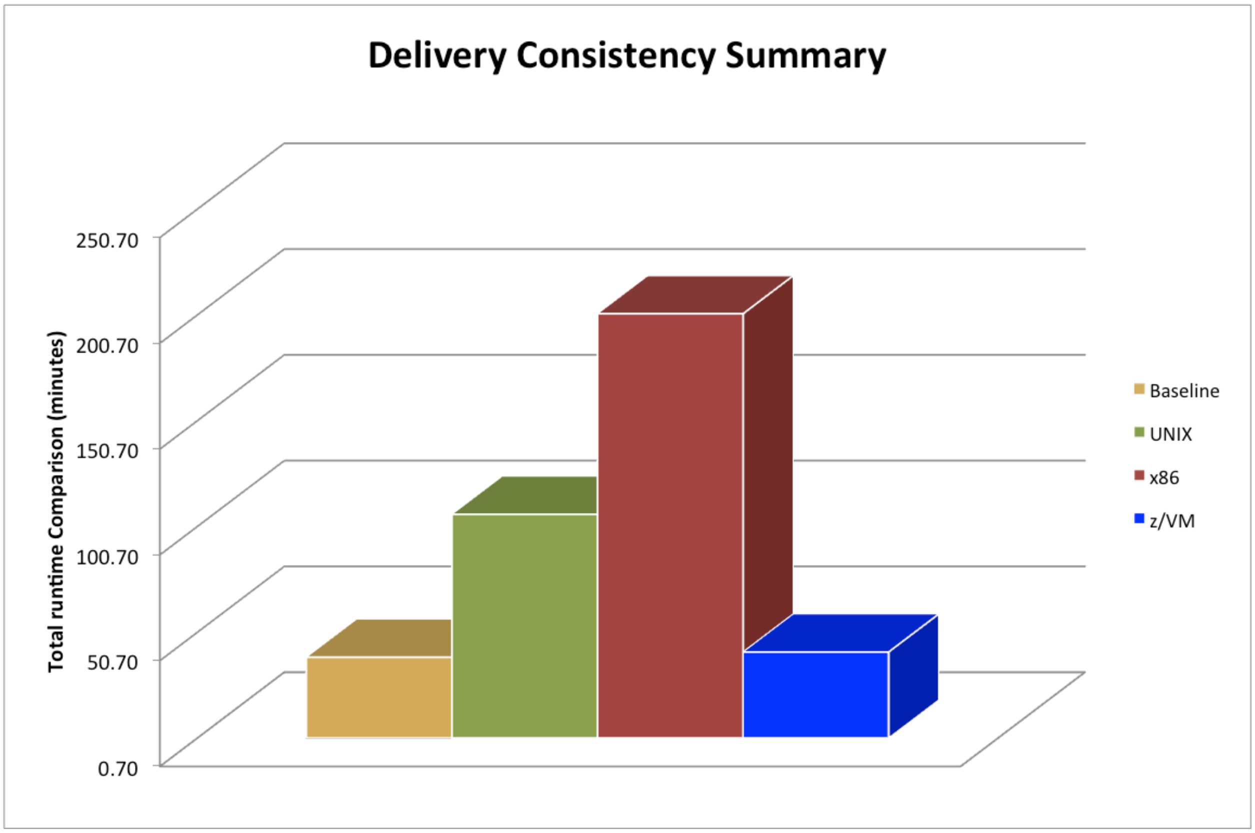

The consistency of the System z platform contributes to the quality of the z/VM deployment, with predictable, reliable services.

System Efficiency

Another component to the technical quality of a virtualization deployment is system efficiency. This is component can be defined as the usage level of the IT asset. The ability to leverage a larger percentage of the IT asset in production is a measure of how effectively the IT asset is utilized by the business. In general, these are consistent within the platform architectural type, with System z averaging 98.8%, UNIX systems averaging 58.6% and x86 platforms averaging 47.5%. This utilization includes the processing power, network load, etc. It should be noted that this is another area that significant variances in the different UNIX and x86 product performance occurs. For specifics on individual offerings, please refer to the SIL study on those architectural products.

Viewed from a virtualization-specific perspective, a key measure of system efficiency is the density of deployed production server VM images. This summary is shown in the graph below. This summary focuses on server images, rather than the more ephemeral and smaller footprint client VMs.

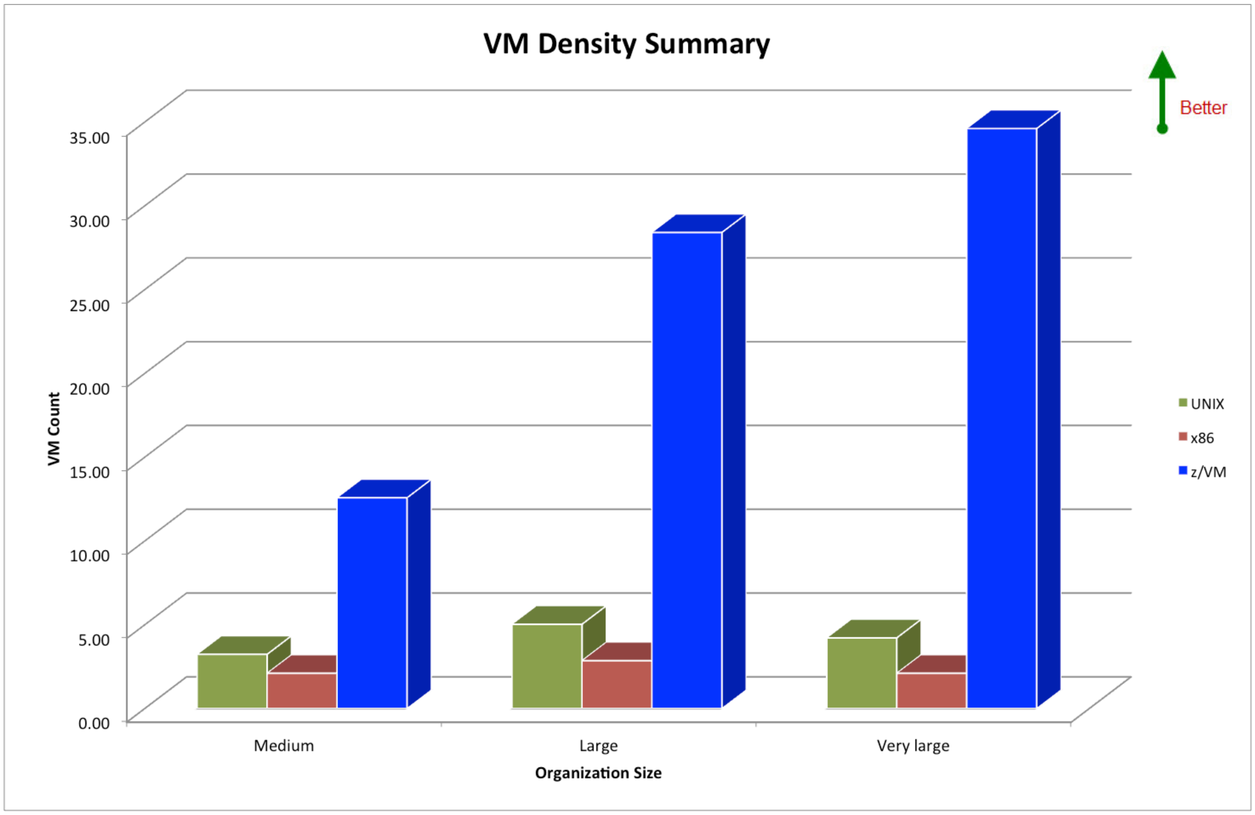

The density of the z/VM deployments is as much as 17 *times* the other platforms. When the reasons for setting the VM density were analyzed, the top three were:

1. Acceptable risk levels
2. Platform performance constraints
3. Organization politics

While the organizational politics are not really pertinent as a technical reason, the first two are especially important for the consideration of a virtualization choice.

The resource utilization per VM provides an interesting view into the technical considerations in this area. The ability of the virtualization method to move resources from one VM to another also comes into play. Effective sharing of resources intra-VM allows the virtualization method to achieve higher levels of overall utilization and load.

In comparisons for this type of metric, the average system utilization is normalized based on the work executed inside of a VM and the cost of a normalized work unit is derived. The cost of this work is then normalized against the set point of a VMware medium-sized implementation. The aggregated metric is based on a SIL-standard comparison point that incorporates total machine instructions and memory byte time usage.

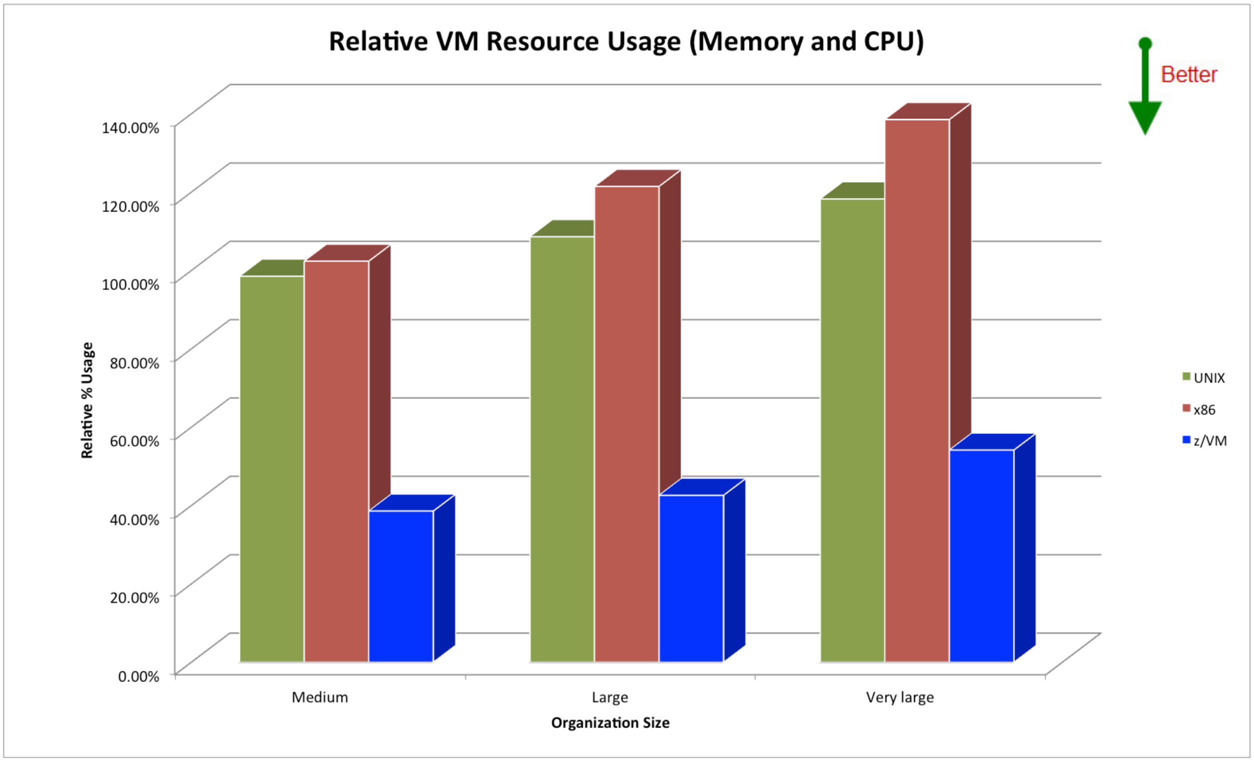

The usage levels of each VM shows clear efficiency for z/VM in resource utilization – an important consideration if fully leveraged platforms are desired. This advantage is as much as 2.55 *times* the competitive offerings.

The cost per work unit for virtualized environments is an indicative metric for those planning a fairly complex environment. In this situation, the resiliency of the underlying architecture is also a substantial contributor to the efficiency of the virtualization methodology.

The cost of deploying each VM is another metric that seems to span both business and technical. These costs include the average cost of platform resource and staff time, but exclude the actual application cost. This is especially important for organizations that have active and volatile non-production environments, since the change in those environments is far higher than that of normal production. When the deployment cost is examined, the comparison is extremely interesting, as can be seen in the chart below. Once again, these costs are normalized against the set point of a medium-sized VMware deployment.

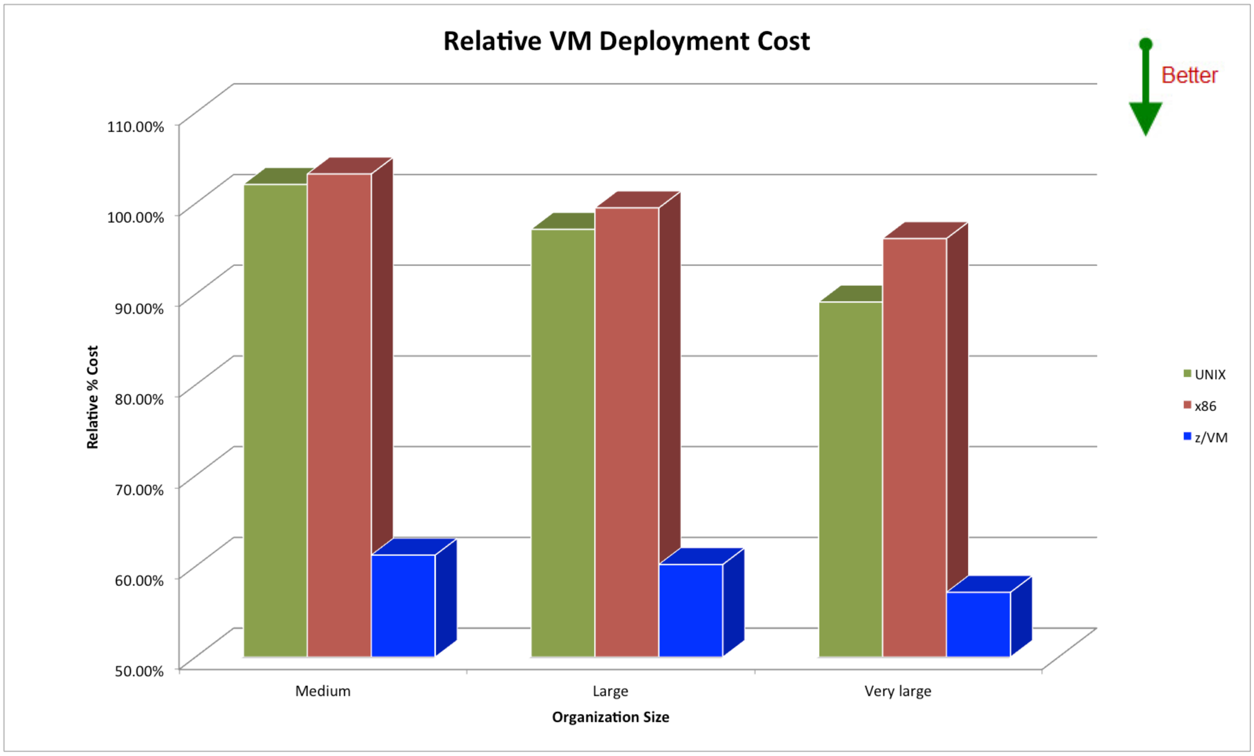

The cost per work unit for virtualized environments is an indicative metric for those planning a fairly complex environment. In this area, the z/VM advantage is as much as 1.68 *times* cheaper than other options. This cost includes hardware, software, staffing and all other normal total expense categories, which is a significant factor for enterprise-wide virtualization considerations.

An underlying limitation that affects all of the technical loads and resource projections for any architecture is the threshold of that platform to tolerate workload. This threshold is the high watermark of sustained usage that still provides a stable and reliable operation. The threshold of sustained use (TSU) is an intrinsic part of the costing and planning when considering where to site IT applications. While there are many different manufacturer claims on operating thresholds, SIL focuses on actual production deployment and loads. Since this reflects the practice, rather than the theory of operation, it provides a more actionable guide for system planning. The following chart summarizes the TSU for the organizations in the study and also provides some insight into underlying risk profiles.

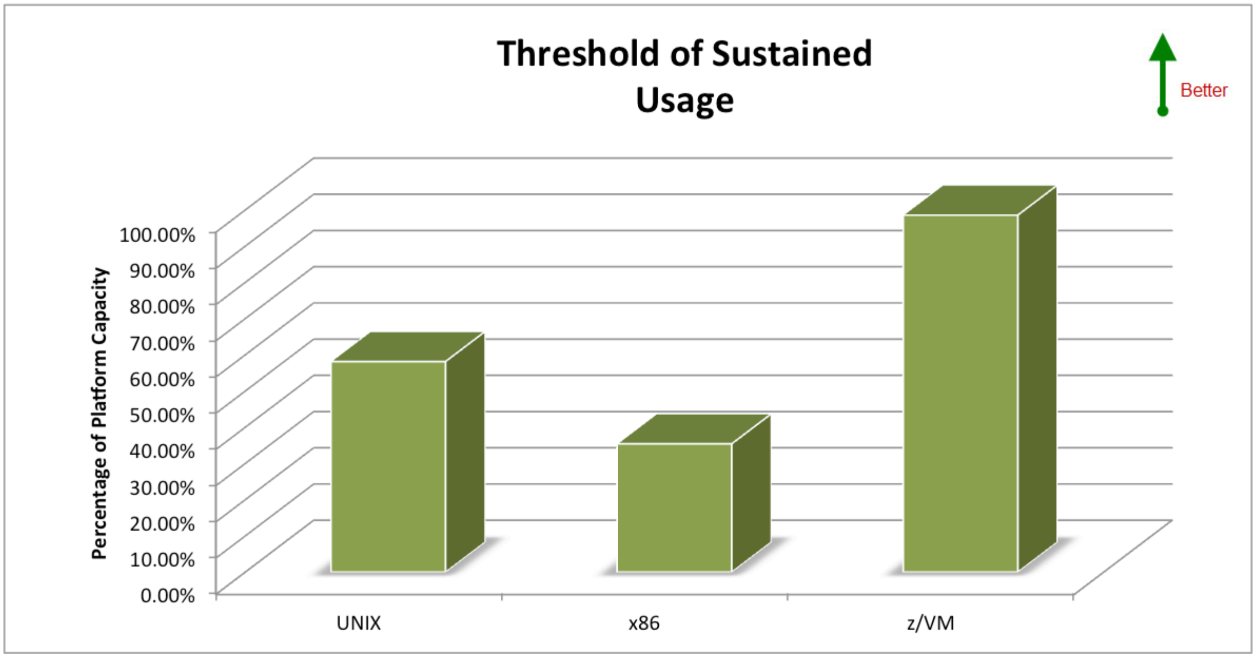

The TSU for z/VM reflects the integrated architecture very clearly, with a sustained threshold that is as much as 2.8 *times* higher than the other competitive offerings.

> **Research Note:** The critical nature of the TSU in creating business cases for architectural selection was clearly demonstrated in the study, which would suggest that this factor be included in an organization's technical selection process.

Environmental Impact

An increasing factor in today's business world is the environmental impact of an organization, including its power and other resource usage, carbon footprint and other eco-impact factors. The most common environmental impacts for measurement are the average floor space, power, cooling and carbon footprint. Although each implementation of each virtualization method can vary widely, the average of those metrics for the different virtualization methods can be used for a high-level indicator for global citizenship. As can be seen below, the System z platform, coupled with the z/VM mechanisms have a synergy that significantly reduces the impact on the environment. This impact affects the square foot area required within a datacenter, the electrical power consumption necessary to run the equipment, the cooling necessary to handle radiated heat within the physical facility, and also the overall carbon footprint. The SIL definition for carbon footprint includes the amount of carbon dioxide produced during the manufacturing process as well as the operational life.

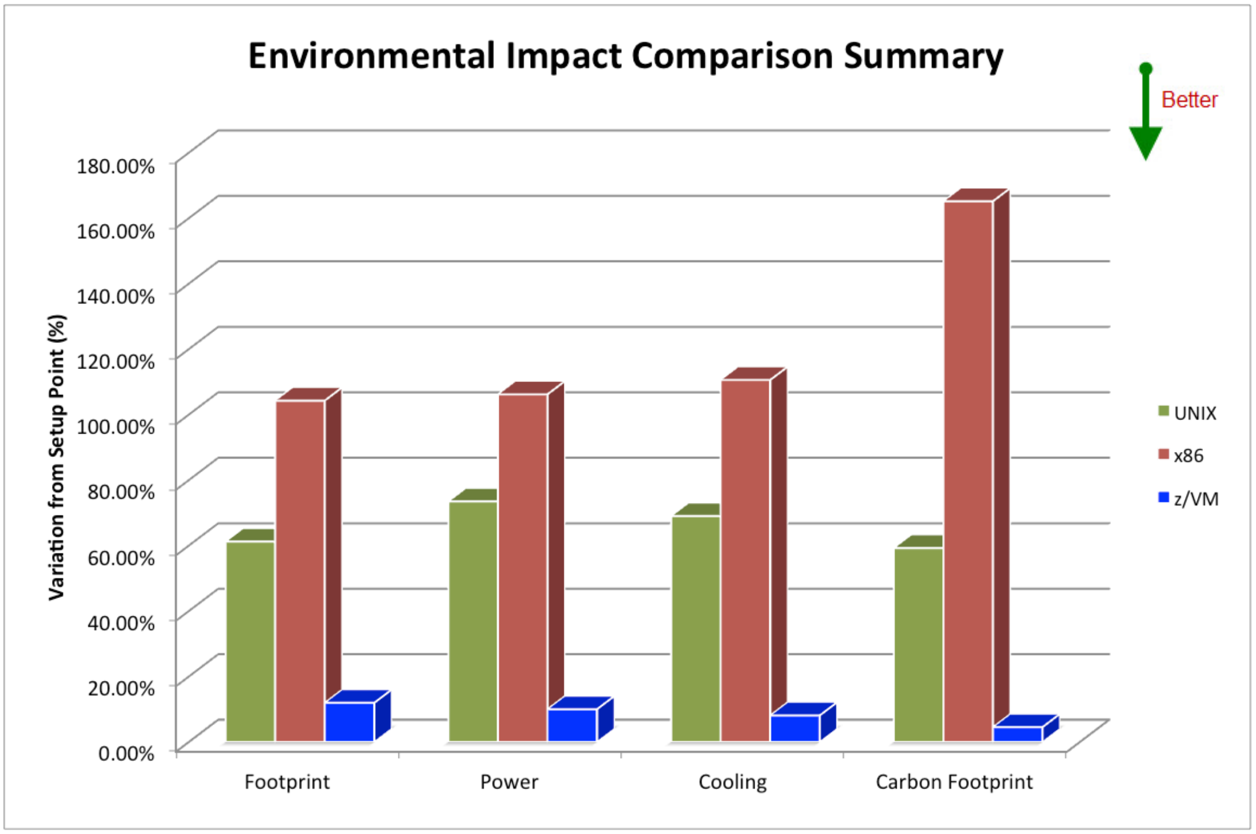

The impact analysis is based on the impact per VM. With the z/VM density of deployment higher levels, the cost to the environment is minimized significantly.

Application Efficiency

Application efficiency is another component of quality that is very seldom addressed, partially due to the difficulty in gathering the information necessary to set this metric. In a simple form, application efficiency reflects how much work is done in an environment to produce the targeted information product. This incorporates the number of machine instructions that are issued to accomplish the different tasks at a machine level, as well as the number of base-level bits that are transferred to and from memory, network, internal storage components, etc. In this area, the underlying OS, coupled with the middleware and the design of the application have to be viewed together, so that the synergies among all of the operating components can be compared equivalently. This affects many of the more external metrics, such as throughput, turnaround, latency, as well as the resource demand of the application.

The view of application efficiency spans not only the integration of the levels of platform architecture, but also the process controls of the actual preparation of an application to run on a specific platform. An example of this is the differences among the platforms to memory leaks. Preparing an application for execution on an x86 platform does not prevent a memory leak from being moved into production, while most UNIX platforms will cavil to greater or lesser degrees. Memory leaks are not tolerated without notification or compile failure on the System z platforms, however. This intolerance has resulted in many applications being corrected for System z execution, which increases the performance and security of the resulting code for that platform. There are many examples of this type of efficiency adjustment, ranging from WAS5, Domino 7, Cognos 8, Oracle 10-11, etc.

This optimization contributes to the application efficiency rating differences to some degree, although the most relevant factor is actually the tight integration of the entire platform stack, including an application that is architected into the System z platform, and incorporated into z/VM. The efficiency also shows up in the number of supportable VMs, the resiliency and other quality factors that differentiate z/VM from the competitive offerings.

The chart below shows the application efficiency rating for the different platforms, averaged across all reporting organizations. It incorporates all of the aspects outlined in this section. A perfect efficiency is considered to be 1.00, so all of these ratings are expressed as a percentage over that perfect one-instruction-one-action level. In other words, the smaller the bar, the more efficient the platform is in each of the areas.

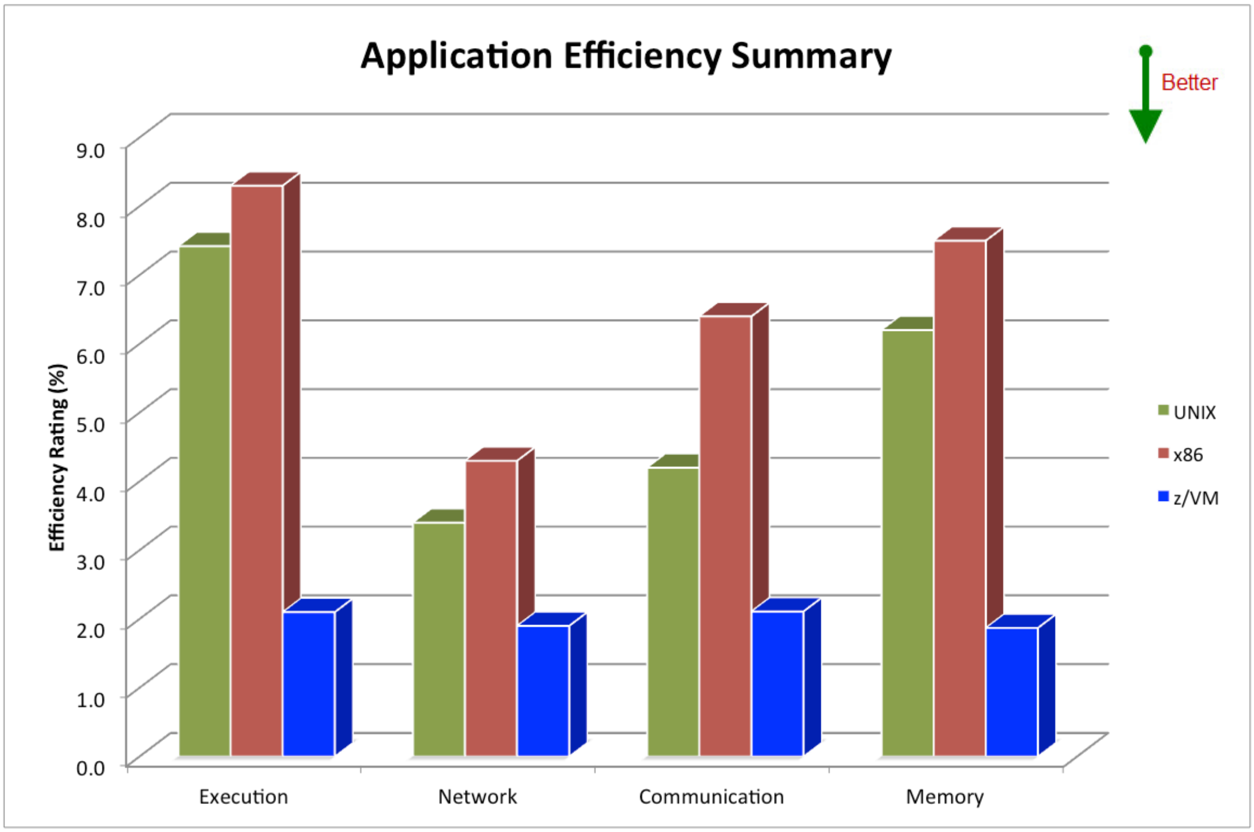

The tightly integrated design of the System z platform spanning hardware, virtualization, operating system and applications creates a level of efficiency that is based on the number of instructions required to execute common compute tasks for delivering business services.

When examining this efficiency in comparison to the competing architectures, the reported experience builds a comprehensive picture. Looking at the extensive base of reporting organizations (79,360+), the average improved efficiency for System z compared to the other architectures in the reporting group can be summarized as:

- application execution - 74.7% better
- network communication - 55.8% better
- communications between applications - 67.0% better
- memory access - 75.1% better

This substantial synergy of System z platform components working with z/VM results in a strong reason to consider z/VM for hosting all applications in a virtualized world.

Management Visibility and Control

The control and management of the environment is also part of the quality matrix. In order to manage IT processing there are some common visibility and control points that can be loosely defined as:

- Monitoring of load and activity
- Data on data – including:
  - events - completion, initiation, etc.
  - coverage - how much data is touched, which jobs are run, where job effects are targeted, etc.
  - timeliness - measured against calendar metric and including the effect on dependencies
  - percentage of data changed
- Data synchronicity - including:
  - data at rest content
  - timestamp of capture
  - correlation pointer - e.g., data warehouse snapshot correlated with its equal point in time of operational data

It is useful to examine the oversight coverage expressed as an average percentage of the production VMs deployed. As can be seen in the following table, there are very spotty coverage percentages for some of the architectures.

Management area	z/VM	x86	UNIX
Monitoring	98.70%	2.10%	6.70%
Metadata	94.30%	1.30%	4.20%
Data synchronicity	74.20%	0.40%	3.60%

Part of the variations in coverage is due to the oversight functions that are present in the native platforms. z/VM automatically has the oversight tools available, while UNIX and x86 has a system activity data collection, with a mechanism for the system activity report, or third-party tools. The cost of equivalent tools on the UNIX and x86 platforms can be extremely expensive, so oversight is often shortchanged on these platforms, which contributes to the higher risk ratings for those platforms.

Security

Security is an important part of any virtualization solution, since virtualization concentrates security topology more densely. The danger to the closely held intellectual capital of an organization is an increasing risk in today's environment.

With the ability to create virtual machines within the same physical platform, the definition of IT security starts to evolve into more than simple access security. The concept of sidewise hacking, where access from one VM to another is broken, like blasting through the walls of an apartment to another within the same building, has started to be a topic of discussion for security personnel everywhere. The protections that the VM software provides have to cover a wider variety of access points than are necessary for security at a whole platform level. In this situation, control over all aspects of processing needs to be in place. Many government and secure installations require protection for the allocation and handling of the main IT spheres: I/O, network access, memory management and overall normal execution access. Within the SIL heuristic database spanning 35+ years of industry oversight, z/VM has no reported incidences of a break in any of the VM security access points.

SIL's view of security is holistic and encompasses a wide scope. It includes:

- Data – access (read, copy) or manipulation
- Process security – ability to execute, hinder, hijack
- Architectural – intellectual property, such as business model, structure of process, metadata
- Physical – access to the physical plant (not in paper purview)

If the security perspective is split into the different classes of security, and a cost analysis is done on what has to be added to the base platform to implement those security levels, an interesting picture is formed. The cost to achieve different levels of security are substantial – sometimes a significant percentage of the overall implementation. To understand these factors, the different security forms can be divided into levels of control:

- Normal corporate
- Credit card processing involved
- Banking
- Healthcare
- Research
- Defense

Based on critical functionality and control, weighted evenly, the different platforms provide the following security coverage, natively.

Security Natively Covered by Platform

Security Level Description	MF	x86	UNIX
Normal corporate	100.00%	21.00%	35.00%
Credit card processing involved	100.00%	14.00%	26.00%
Banking	100.00%	8.00%	14.00%
Healthcare	100.00%	7.50%	11.00%
Research	78.00%	3.00%	8.00%
Defense	64.00%	1.00%	3.00%

If the cost to bring each of the architectures up to the required level of security (100%) is tabulated, it can be expressed as a percentage of the total cost of the implementation.

Incremental Cost to Achieve Required Security

Security Level Description	MF	x86	UNIX
Normal corporate	0.00%	25.20%	12.10%
Credit card processing involved	0.00%	38.40%	16.90%
Banking	0.00%	63.70%	22.40%
Healthcare	0.00%	81.60%	30.70%
Research	2.10%	134.80%	56.90%
Defense	4.30%	187.90%	97.50%

This cost summary includes operations, hardware, middleware, but no applications. It does include a z/VM implementation that has Tivoli Identity Manager, RACF for z/VM and zSecure. The Trusted Key Entry appliance was not included in this base, although it is available for an additional charge.

During separate study activities, SIL has conducted a series of vulnerability analyses for a random group of customers. Some of those customers were aware of security incursions of one sort or another, but all were concerned with a view into the effectiveness of their security. A total of 599 customers were analyzed in detail during the SIL Vulnerability studies out of the total customers present in this main study. The most surprising finding that came from this targeted analysis is that the large majority of those customers were not aware of the actual incursions into their systems. In fact, in many cases optimizations that were put into place to address performance or application challenges, such as slow response, had created pockets of vulnerability within the IT environment, such as caching passwords in the clear on internal networks. In general, some of the organizations were aware of security breaches, but some surprising findings were observed. The table below summarizes these results.

Description	Count	Security Summary A45023 – August 18, 2012			
Customers in security and vulnerability analysis	599				
Incursion Description	**Total Incursion**	**Incursion %**	**Mainframe**	**x86**	**UNIX**
Aware of security breach	41	6.84%	2	0	39
Organizations with security breaches	565	94.32%	2	512	169
Peek at data	565	94.32%	2	491	237
Copy data	511	85.31%	0	503	221
Change data	168	28.05%	0	154	39
Affect process	305	50.92%	0	292	194
Extract metadata (IP)	117	19.53%	0	115	29
Employee security breaches	68	11.35%	2	68	16

It should be noted that the total number of breaches across architectures might exceed the number of base incursions. This occurs when an organization has security breaches in more than one architecture.

The most startling finding was the sheer number of organizations that had experienced security breaches of which they were unaware. During this random set of vulnerability checks, 42 organizations had alien extraction processes that were still in active piracy mode, stealing the information and affecting processes in real time. It was notable that the only two security breaches that occurred in z/VM systems were from employee misuse of internal security protocols.

With z/VM, the level of execution, process and I/O isolation allows the applications from multiple competing ISVs to be supported while co-located within a single virtual image. What this means is that z/VM is the only one of the architectures reviewed in which competing ISVs officially support their applications in a shared virtual environment, e.g., DB2 and Oracle DB in the same space. This is a result of the synergistic architectural security design and contributes to the significant difference in the risk associated with the System z and z/VM deployments.

> **Research Note:** The lack of awareness concerning critical security breaches would suggest that a more aggressive monitoring and access protocol be investigated for most organizations. This should minimally include all the forms of security breach that have been listed.

Cloud Integration

The integration with cloud services is a very common initiative in today's marketplace. With that in mind, part of the study looked at organization deployment on the cloud and the link to an organization's virtualization strategy. In this area there were several points of analysis. The first of these is the customer use of cloud, split out by virtualization method.

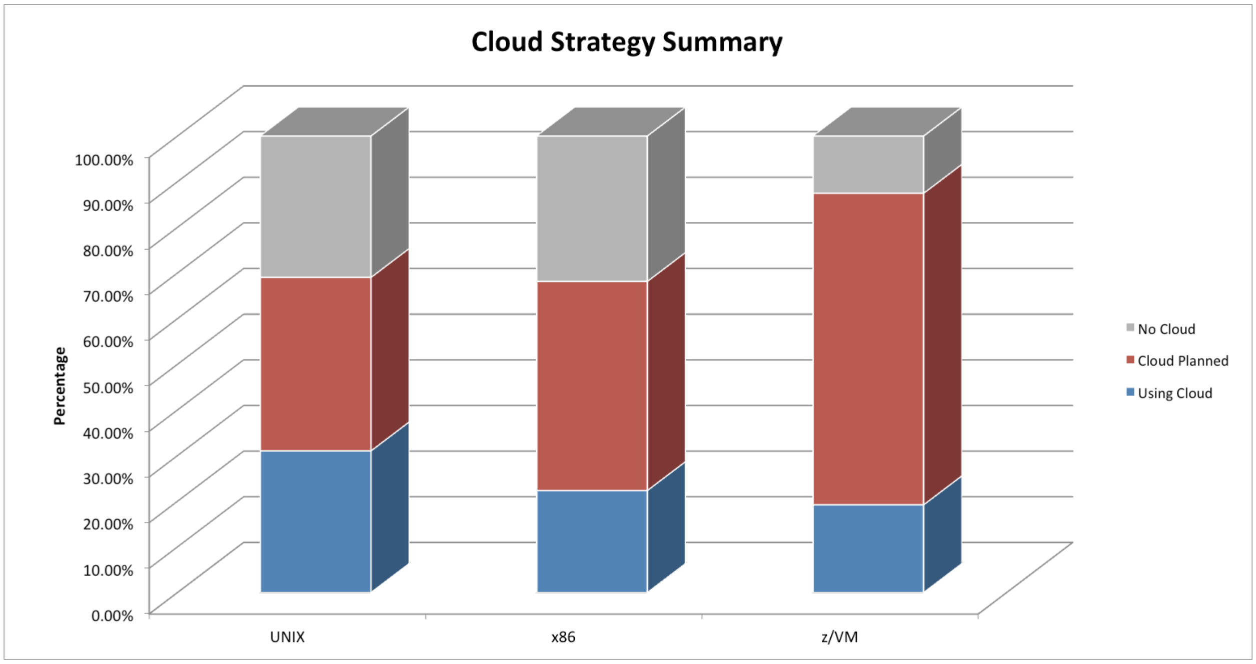

The second perspective is which customers have selected their virtualization with cloud deployment (either current or planned) in mind. Of those responding, more than 81% said that cloud was a consideration in selecting their virtualization method. Both of these perspectives have created an interesting viewpoint into the cloud movement. Since cloud architecture is really a further form of virtualization, the selection of the architectural strategy says a lot about which technology has the robustness to position an organization for the future. To better understand the trends of cloud integration, an adoption trend was done for each of the main architectural choices.

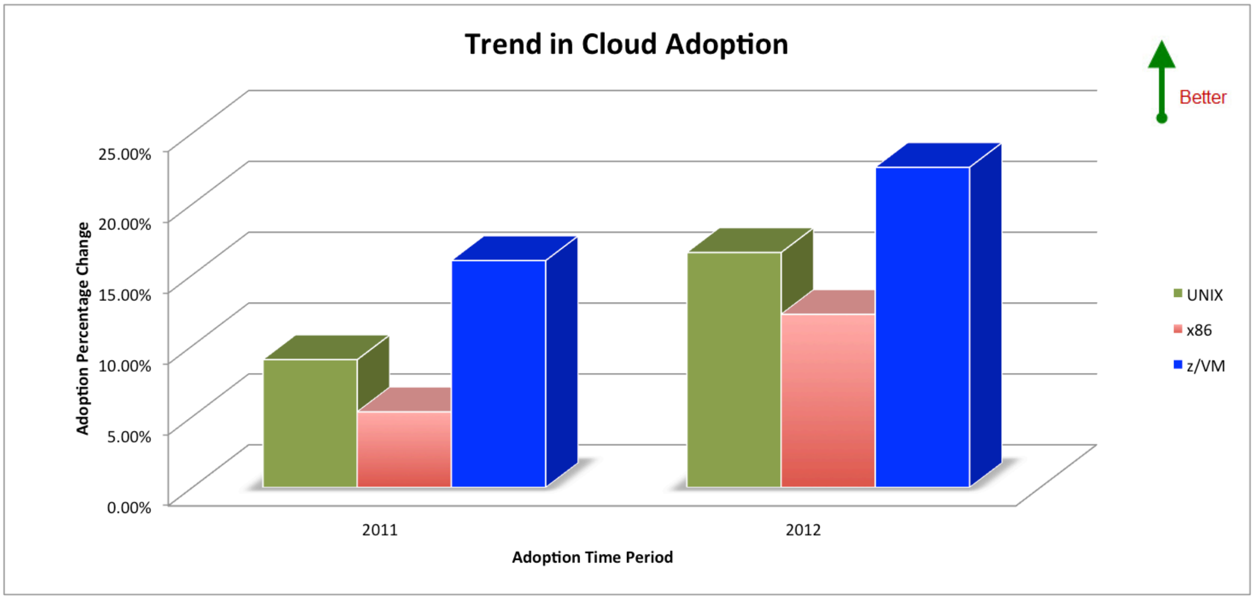

As can be seen above, the integration of clouds using z/VM is growing rapidly. Even with the trend limited to the last 18-month period, the adoption of z/VM to support private and public clouds is impressive, growing by 48.6% over the study period.

2.5. Platform Relevancy

Over the years, the entire mainframe architecture has fallen into a perception of overkill. Many organizations have moved away from the System z architecture, viewing it as too expensive, too difficult and too outdated for consideration. However, a real understanding of TCO and other factors is changing this perception in the marketplace. The main influences pushing the change in perception are:

- Better tracking of IT TCO
- Expansion into multi-time zone (6+) markets
- Increased need for uptime and availability
- Security concerns
- Staffing reductions
- Rapid market changes, requiring faster implementation response
- Increased need for application resiliency

All of these increasing requirements on the enterprise level have made the business case for mainframe deployment more and more relevant. SIL has tracked architectural relevance for a considerable number of years, broken down into three categories:

- Business case justified
- Possible justification
- Not justified

These categories are simplistic and are summaries of clear classification, so some of the "possibles" may actually be justified when additional factors are included. However, for the purposes of this summation, the more conservative categorization is appropriate. The graph below shows the split among these three categories for all organizations within the study based on known workload, cost, availability, performance and complexity requirements. The relevancy snapshots used are August 2010, August 2011 and August 2012. A final snapshot is the result of projective modeling, with all known initiatives reflected in the projection of workload and other requirements for the upcoming year.

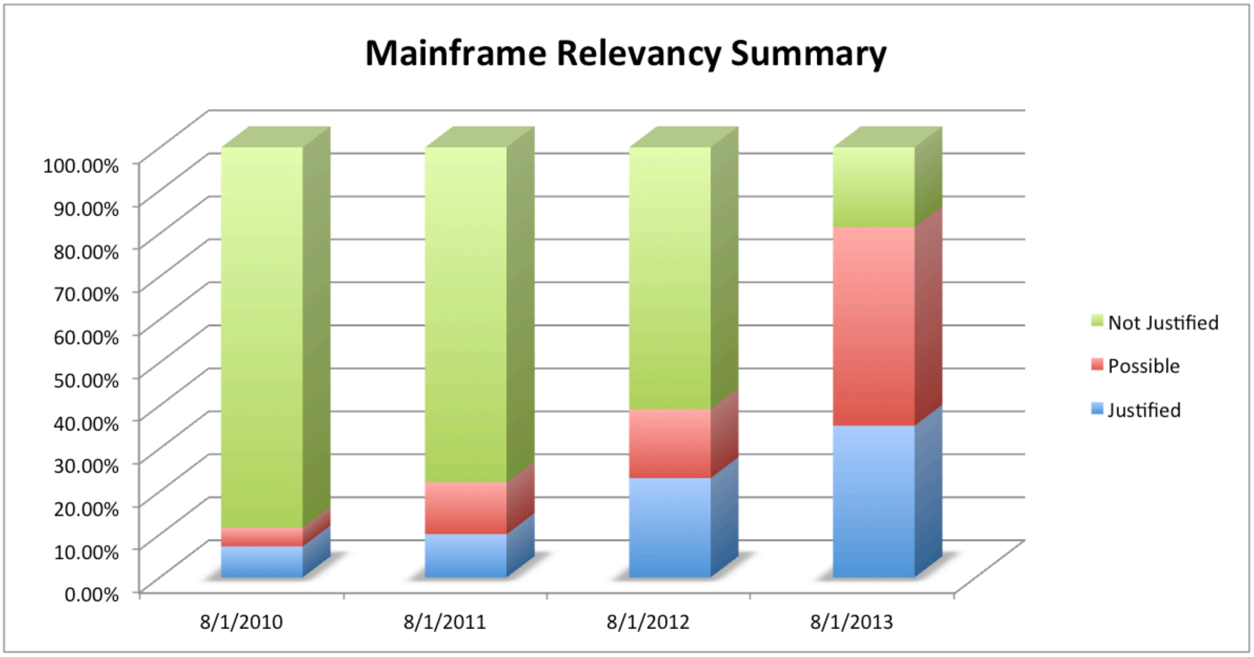

The increase in the use of business analytics, deployments across time zones, global operations and other factors have increased the complexity, raised the workload and pushed availability limits for many organizations. Where an organization has realized its need for doing business on a reliable basis across time zones, maintaining strong data synchronicity, handling big data or data analytics, or even increasing its awareness of security breaches, the mainframe architecture has increased its relevance considerably.

> **Research Note:** The massive changes in requirements for organizations using extensive business analytics, big data sources or with stringent security and availability requirements suggests that additional strategic evaluation be done to accurately reflect modern architectures with modern workload.

2.6. Conclusion

The Solitaire Interglobal Ltd. analysis of operating systems shows that the IBM z/VM virtualization method has a significant set of benefits for a wide range of organizations, based on a broad set of business and performance metrics. The advantages that accompany this platform increase the effectiveness of application deployment and translate to real-world positive results experienced and reported by the businesses in this study. In fact, more than 80% of the study's organizations have a strong business case for z/VM deployment, when viewed with cost, workload, availability, security and IT services quality in mind.

While success can be measured in different ways and looked at from varying perspectives, it could be said that the bottom-line measurement of deployment success is overall customer satisfaction. Customer satisfaction incorporates a wide variety of qualitative and quantitative components, yet it is the simplest summary of how well a deployed system has met organizational expectations. As outlined in the analysis, the customer satisfaction with the z/VM choice is high in more complex environments, both from a business and from a technical perspective.

The economic benefits of the virtualization choice are also apparent in the control of overall expense. This study has identified critical business and performance metrics that can be used to understand the advantages and key strategies that will help an organization to choose the optimal operating system.

Whether an organization is looking to deploy cloud architecture, or simply desires a quality, consistent, dependable and efficient IT deployment in a virtualized world, the choice of virtualization strategy is critical. Any organization considering all of the factors will see that z/VM is a relevant option worthy of consideration.

About Solitaire Interglobal Ltd.

Solitaire Interglobal Ltd. (SIL) is an expert services provider that specializes in applied predictive performance modeling. Established in 1978, SIL leverages extensive AI technology and proprietary chaos mathematics to analyze prophetic or forensic scenarios. SIL analysis provides over 4,500 customers worldwide with ongoing risk profiling, performance root cause analysis, environmental impact, capacity management, market trending, defect analysis, application Fourdham efficiency analysis, organizational dynamic leverage identification, as well as cost and expense dissection. SIL also provides RFP certification for vendor responses to government organizations around the world and many commercial firms.

A wide range of commercial and governmental hardware and software providers work with SIL to obtain certification for the performance capabilities and limitations of their offerings. SIL also works with these vendors to improve throughput and scalability for customer deployments and to provide risk profiles and other risk mitigation strategies. SIL has been involved deeply in the establishment of industrial standards and performance certification for the last several decades and has been conducting active information gathering for the Operational Characterization Master Study (OPMS) – chartered to develop better understanding of IT-centric organizational costs and behavioral characteristics. The OPMS has continued to build SIL's heuristic database, currently exceeding 95 PB of information. The increased statistical base has continued to improve SIL accuracy and analytical turnaround to unmatched levels in the industry. Overall, SIL runs over 40,000 models per year in support of both ongoing subscription customers and ad hoc inquiries.

Further Methodology Notes

In order to support the comprehensive nature of this analysis, information from diverse deployments, industries, geographies, and vendors were obtained. In any collection of this type, there is some overlap that occurs, such as when multiple vendors are present at an organization. In such cases, the total of the discrete percentages may exceed 100%. Those organizations with a multi-layered deployment, such as multiple geographical locations or industrial classifications, have been analyzed with discrete breakouts of their feedback for all metrics. Additional filtering was performed to eliminate those implementations that substantially failed to meet best practices. Since the failure rates, poor performance and high costs that appear in a large number of those implementations have little to do with the actual hardware and software choices, these projects were removed from the analytical base of this study.

The industry representation covers manufacturing (29.62%), distribution (17.89%), healthcare (10.86%), retail (5.39%), financial (11.70%), public sector (12.93%), communications (7.99%) and a miscellaneous group (3.61%).

The geographies are also well represented with North and South America providing 44.88% of the reporting organizations, Europe 28.53%, Pacific Rim and Asia 22.72%, Africa 3.33%, and those organizations that do not fit into those geographic divisions reporting 0.54% of the information.

Since strategies and benefits tend to vary by organization size, SIL further groups the organizations by the categories of small, medium, large and extra large. These categories combine the number of employees and the gross annual revenue of the organization. This staff count multiplied by gross revenue creates a metric for definition that is used throughout the analysis. In this definition, a small organization could be expected to have fewer than 100 employees and gross less than $20 million, or a value of 2,000, e.g., 100 (employees) X 20 (million dollars of gross revenue). An organization with 50 employees and gross revenue of $40 million would have the same size rating, and would be grouped in the analysis with the first company. The classifications used by SIL use thresholds of 2,000 (small), 10,000 (medium), 100,000 (large) and 1,000,000 (extra large).

Attributions and Disclaimers

IBM, z/VM, zLinux and System z are trademarks or registered trademarks of International Business Machines Corporation in the United States of America and other countries.

UNIX is a registered trademark in the United States of America and other countries, licensed exclusively through The Open Group.

Other company, product and service names may be trademarks or service marks of others

This document was developed with IBM funding. Although the document may utilize publicly available material from various vendors, including IBM, it does not necessarily reflect the positions of such vendors on the issues addressed in this document.