ITG Executive Summary

November 2008

VALUE PROPOSITION FOR IBM SYSTEMS DIRECTOR: CHALLENGES OF OPERATIONAL MANAGEMENT FOR ENTERPRISE SERVER INSTALLATIONS

Economics Benefits for IBM System z Deployment

Challenges

At the end of 2003, the typical U.S. Fortune 1000 corporation contained fewer than 3,000 servers. On current trends, by the end of 2008 this will have increased to more than 4,800. The vast majority of these – almost 88 percent – will be x86 platforms running Windows or Linux operating systems.

The challenges created by physical server proliferation have been well documented. They include low overall capacity utilization, high management overhead, difficulties in maintaining service quality, escalating energy costs, and – in many data centers – pressures to add space, upgrade energy and cooling infrastructures, or both.

Growing adoption of x86 server virtualization has begun to slow growth in physical server bases. But, as organizations have moved beyond pilot projects and begun widespread deployments using such enablers as VMware, Microsoft Virtual Server, and Xen, they have discovered new sets of challenges.

A new layer of software is added to enterprise x86 environments. Virtual as well as physical server resources must be managed. Maintenance of availability, backup and recovery, and security becomes more problematic when multiple virtual servers run on single physical platform.

A new trend must also be dealt with – virtual server proliferation. Numbers of virtual images are expanding rapidly. If current trends continue, by the end of 2013 the typical Fortune 1000 corporation will contain more than 6,500 physical and between 5,000 and 10,000 virtual servers. Managing such environments will be a significantly more demanding task than managing physical server bases.

Server virtualization offers important benefits. It is becoming clear, however, that popular x86 server virtualization enablers offer an incomplete solution. Alternatives that offer more mature and efficient virtualization capabilities should also be reviewed.

This report deals with one such alternative – the use of the IBM System z10 platform with Linux guests and the new IBM Systems Director 6.1 platform management solution.

In six composite profile installations of IBM System z10 Enterprise Class (EC) and Business Class (BC) systems, replacement of distributed Windows, Linux, and Solaris servers with z10-attached Integrated Facilities for Linux (IFLs) enabled major savings in personnel costs for server administration. Three-year costs were reduced by an average of 38.1 percent.

If IBM Systems Director 6.1 was also employed, additional server administration savings were realized. Costs were reduced by an average of 47.9 percent.

Solutions

The System z platform has been used for server consolidation for almost a decade. More than 25 percent of System z customers have installed IFLs to support, and more than 2,500 applications have been written or migrated to run on SUSE, Red Hat, and other Linux distributions on System z.

To date, most System z server consolidation activity has been among large IBM corporate and government customers. The new System z10 Business Class, however, is positioned as a consolidation solution for smaller organizations. Linux performance is said to be up to 50 percent better than for the earlier z9 Business Class, and pricing for Linux workloads is more aggressive than for the latter.

The value proposition for the System z10 for server consolidation is reinforced by Systems Director 6.1 support for System z Linux environments. Systems Director 6.1 replaces the earlier IBM Director.

Systems Director 6.1 provides integrated discovery, monitoring, configuration and event management, software update, and other services for the full range of IBM server platforms. IBM AIX and i, Windows, and Linux operating systems are supported, along with IBM PowerVM (for Power servers), VMware, Microsoft Virtual Server, and Xen virtualization enablers.

Management services are also extended to IBM distributed disk systems; SAN Volume Controller (SVC) storage virtualization software; storage area network (SAN) switches and network devices; and other IBM and non-IBM systems and devices. A common management console is provided for all of these.

This approach, which is illustrated in figure 1, provides a unified management capability for a wide range of IBM platforms. Certain non-IBM platforms and SNMP devices such as printers and network devices are also supported.

Figure 1
Systems Director 6.1 Cross-platform Management Capability

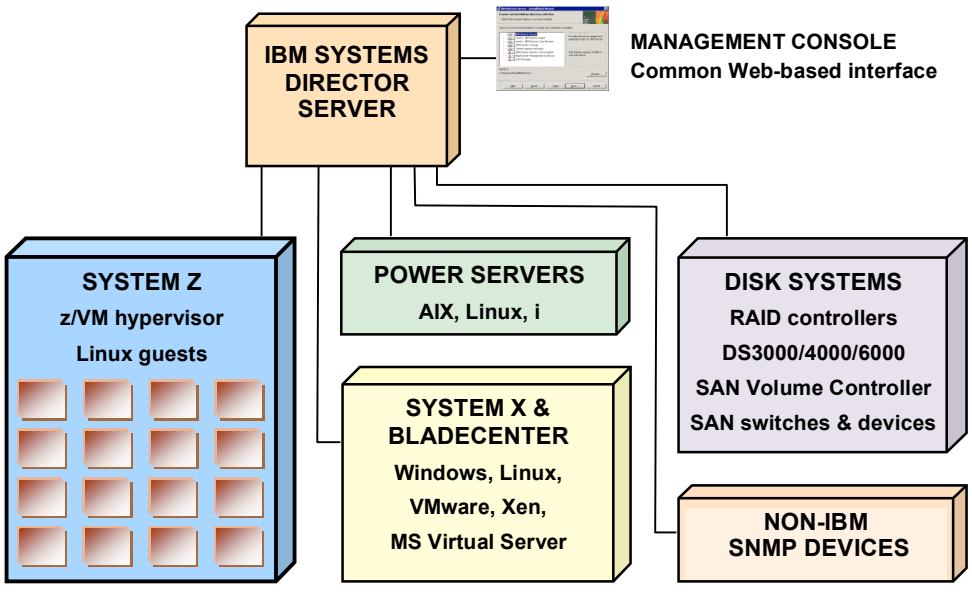

The combination of System z and Systems Director 6.1 delivers dual benefits: it enables new management efficiencies within the System z Linux environment; and it allows organizations to standardize server administration tasks for System z and other IBM platforms and to employ a single browser-based administrator interface for these.

The cost savings potential of both effects is significant.

Savings

The six installations for which costs were compared for this report included banking, distribution, financial services, insurance, and retail companies with between $120 million and $20 billion in sales and 1,000 to 35,000 employees, along with a state government agency with approximately 10,000 employees.

Profile installations were based on data supplied by 13 organizations in the same industries and approximate size ranges. These organizations had used System z IFLs for server consolidation initiatives. System z10 platforms with between one and six IFLs were employed to consolidate between 45 and 450 distributed servers.

Three-year server administration costs were compared for these installations for three scenarios: (1) replaced servers; (2) use of System z Linux without Systems Director 6.1; and (3) use of System z Linux with Systems Director 6.1.

When Systems Director 6.1 was employed, System z Linux costs ranged from 26.9 percent to 53.5 percent less.

Three-year savings were higher for installations with four to six IFLs – combined use of System z Linux and Systems Director 6.1 resulted in server administration costs that averaged 49.4 percent less than for replaced servers. For installations with one to three IFLs, however, savings were also substantial – costs averaged 41.9 percent less than those for replaced servers.

These results are summarized in figures 2 and 3.

Figure 2

Three-year Server Administration Costs for Server Scenarios: Installations with 4 – 6 IFLs

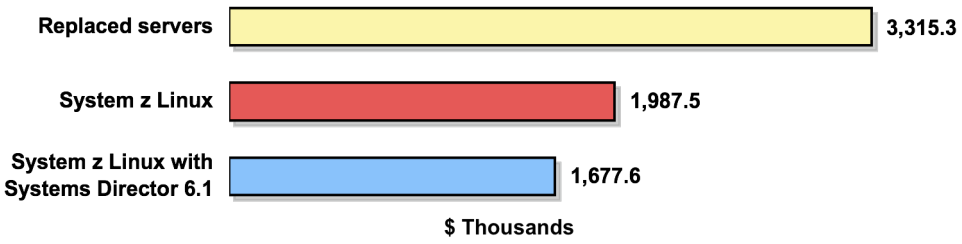

Source: International Technology Group

Figure 3

Three-year Server Administration Costs for Server Scenarios: Installations with 1 – 3 IFLs

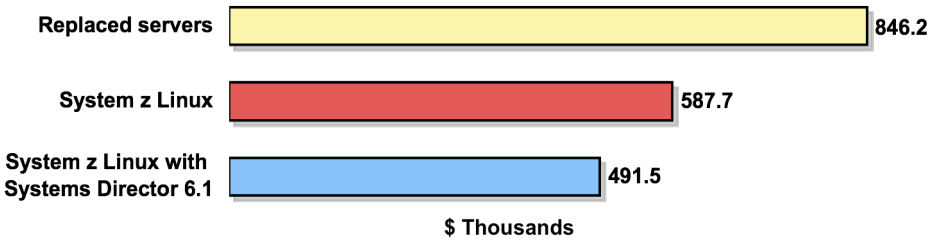

Source: International Technology Group

Higher savings for installations with from four to six IFLs reflect two factors: (1) more servers were replaced by System z Linux virtual servers than in installations with one to three IFLs; and (2) Systems Director 6.1 could be employed to manage larger numbers of IBM platforms of all types.

Additional information on installations, along with details of methodology and assumptions, may be found in the Basis of Calculations section of this report.

A related IBM product, IBM Systems Director Active Energy Manager (AEM), can also be used through the Systems Director 6.1 management console. AEM, which can run in a System z Linux partition, can be used to monitor energy usage by and thermal loading for System z as well as other IBM platforms. This capability has formed an important part of data center energy conservation programs in many organizations.

Positioning System z

User View

Most organizations that contributed data to this report had conducted, or were in the process of conducting server consolidation initiatives using multiple techniques.

In addition to use of System z Linux, these included conventional consolidation (i.e., replacement of distributed servers running the same application with fewer, centralized platforms) and use of VMware, PowerVM, and equivalents to consolidate heterogeneous applications and server functions onto shared x86 or Power platforms.

One organization, for example, planned to reduce its server base from around 4,500 Windows, Linux, and UNIX physical servers to approximately 1,800 using a combination of these techniques. Another planned to replace more than 600 distributed Windows and Linux servers using VMware-equipped System x and BladeCenter platforms and System z IFLs.

In these and other cases, it was reported that the System z platform offered distinctive advantages in two areas – technical capabilities and economics – that may be summarized as follows.

Technical Capabilities

The following technical capabilities were cited by users:

- ***Virtualization***. System z virtualization capabilities were variously characterized as "mature," "robust," "stable," and "solid."

  The System z platform offers two approaches to hosting Linux guests: (1) the hypervisor-based z/VM operating system and (2) firmware-based logical partitions (LPARs). Many organizations employ both.

  These are the IT world's most highly developed and mature virtualization enablers. z/VM is derived from the IBM VM operating system, which was introduced in 1972, while LPAR capability was introduced in 1988. Linux deployments using these began in the late 1990s.

  Both enablers are implemented in a highly efficient manner – there is virtually no emulation overhead – and are integrated with industry-leading capabilities for virtualization of processors, memory, I/O, and disks. z/VM also supports high-speed virtual LANs that handle communications between virtual servers within the system.

This is a key differentiator. The effectiveness of a virtualization solution is determined not only by its hypervisor. Overall performance and efficiency are also affected by other system components. When a full range of virtualization technologies are integrated and optimized – which is the case for the System z platform – the whole may be far greater than the sum of the parts.

- ***Scalability***. There was general agreement that larger numbers of virtual images could be consolidated on System z than on competitive platforms.

  The System z has a long history of scalability. Some organizations have employed this platform and its predecessors to handle sustained growth over decades. The latest z10 generation of systems supports up to 64-way configurations with very low levels of symmetric multiprocessing (SMP) overhead.

  Up to 10 IFLs are supported on z10 BC and up to 64 on z10 EC models. Assuming a typical range of approximately 15 to 150 Linux images per IFL, the z10 BC can support 150 to 1,500 Linux virtual servers, and the z10 EC can support approximately 1,000 to 10,000.

  A further level of capability is provided by Parallel Sysplex clusters, which can enable clusters of up to 32 nodes for a maximum of 64 x 32 = 2,048 processors. Working clusters in the 4- to 8-node range are routine, and some are in the 10- to 20-node range.

- ***Availability***. This was cited as a key System z strength. Several organizations had decided to consolidate servers supporting business-critical applications to System z while employing x86 virtualization enablers for others that were less sensitive to downtime.

  The System z is generally recognized as the IT world's most reliable server. Most users experience no unplanned downtime over multi-year periods – in some cases, for more than a decade. Planned outages are also typically rare and short.

  Availability levels of 99.99 to 99.999 percent are common, and some organizations experience 100 percent availability. (This is more common than is generally realized, at least for mainframe systems. IT managers are often reluctant to commit to "100 percent" availability in service level agreements, even if this is realized in practice.)

  The System z benefits from extremely high levels of component redundancy and from industry-leading embedded diagnostic, fault prevention, and fault resolution mechanisms. All major components, including processors, can be replaced or serviced without taking systems offline.

  IFLs, which are physically the same as control processors (CPs) employed for mainstream z/OS workloads, benefit from all System z hardware-based availability features. z/VM is highly stable and equipped with extensive resiliency capabilities. Software upgrades and maintenance actions may be performed in LPARs without disrupting system operations.

- ***Workload management***. Server consolidation requires more than the creation of virtual images.

  Virtualization creates the potential for significantly higher levels of overall capacity utilization than may be realized with dedicated servers. The extent to which this potential will be realized in practice, however, depends heavily on the mechanisms that allocate resources between, and monitor and control workload execution processes across virtual servers.

  If these mechanisms are ineffective, a high proportion of system capacity may be idle at any given time. Surges in workloads for individual virtual servers may also create bottlenecks if additional capacity is not available in a timely manner.

The System z platform is a recognized industry leader in workload management. It is equipped with mechanisms that allow system resources to be allocated and dynamically reallocated with extremely high levels of granularity based on workload priorities and on response time, job execution time, and other service level agreement (SLA) variables.

Although the highest levels of workload management granularity are supported for the z/OS environment, System z Linux guests benefit from a level of capability that is significantly better than for most competitive platforms.

- ***Security.*** The ability to maintain high levels of security in a comparatively simple and cost-effective manner was also seen as a key System z advantage.

  In addition to industry-leading physical security, extensive security mechanisms are built into z/VM. While this does not directly secure Linux guests, it makes z/VM less vulnerable to exploits than, for example, VMware, Microsoft Virtual Server, or Xen.

  Additional security functions are built into LPAR firmware. z/VM virtual LANs and HiperSockets links, which eliminate external network connections, further reduce vulnerabilities. Established System z security tools, such as the IBM Resource Access Control Facility (RACF), may also be leveraged.

  System z Linux guests are supported by the Crypto Express2 feature, an onboard encryption co-processor that implements Secure Sockets Layer (SSL) and other industry standards including Federal Information Processing Standard (FIPS) 140-2 Level 4. This approach is more efficient than use of outboard engines.

- ***Disaster recovery***. Server consolidation increases vulnerability to disruptions. If a single platform hosts dozens or even hundreds of virtual images, large numbers of applications and user groups may be impacted by a system or data center failure affecting it.

  System z Parallel Sysplex clustering enables rapid and reliable local failover. The company's Geographically Dispersed Parallel Sysplex (GDPS) solution, an extension of Parallel Sysplex, is recognized as the world's premier solution for wide area failover and recovery to remote data centers. Both provide full coverage of System z Linux virtual servers.

  GDPS enables near real-time failover at distances of up to 300 kilometers. Distances of 30 to 100 kilometers are routine among users, who include most large financial institutions and electronic funds transfer networks, as well as large corporations in a wide range of other industries and major government organizations worldwide.

  The experiences of these users have been consistent with the IBM position that, for large operational sites, a planned switch from one site to the other requires less than 60 minutes (including network delays) and site unplanned outage recovery takes less than 45 minutes.

Similar capabilities, at least to some extent, may be realized for distributed servers. For example, availability levels and disaster recovery may be improved by clustering Windows, Linux, and UNIX platforms. These may also be equipped with workload management, security, and other tools that approximate System z functionality.

These approaches, however, are complex and expensive if they are applied on an application-by-application, server-by-server basis. Most organizations that employ them restrict their use to the most business-critical x86 or UNIX server-based systems.

In contrast, consolidation to System z enables organizations to extend business-critical capabilities to all virtual servers running on this platform with minimal difficulty, at no additional cost.

According to users, however, this was not the only economic benefit of employing System z IFLs.

Economics

The System z platform, when running the z/OS operating system and equipped with comparatively expensive suites of IBM and third-party software supporting core business-critical systems, is not generally regarded as a "low-cost" platform.

The IBM pricing model for System z IFLs, however, is significantly different. Organizations that contributed to this report regarded IFL cost structures as competitive with those of x86 or UNIX servers.

One reason for this was that System z-based data center infrastructures and skill bases were already in place and were cost-justified by core business systems. Existing System z platforms were already supported by disk and tape storage systems, networks and communications facilities, and power and cooling infrastructures that could be expanded at little or no cost to handle new Linux workloads.

In most organizations, existing backup and recovery mechanisms, Parallel Sysplex and GDPS clusters, and operations support systems and processes could also be expanded to cover Linux IFLs.

Equally, while it might be necessary to add or transfer staff for Linux development and administration, existing System z specialists in system and storage management, backup and recovery, operations, and other data disciplines could typically handle these tasks for IFLs as well as for z/OS systems.

Users also cited the advantages of IBM Capacity Upgrade on Demand (CUoD) arrangements, which enabled organizations to install IFLs and to pay for these only when they were activated. These were seen as particularly valuable by organizations that were experiencing sustained workload growth, or were subject to periodic workload spikes, or both.

One retailer, for example, employed additional IFLs to handle increased workloads during the Christmas shopping season. Another noted that the company was able to manage a spike in Web site activity during a major promotional campaign by "renting an IFL for two weeks."

Enterprise View

For more than a decade, there has been a growing recognition that all components of organizational IT environments have grown more complex and interdependent. It has become clear that full range of IT resources must be managed more effectively at the enterprise level.

The challenges, however, are daunting. Enterprise management capabilities must extend across servers and storage as well as networks, applications, and middleware. They must integrate traditional disciplines such as systems management, operations, asset management, change management, and service-level management, along with increasingly critical security, compliance, and energy conservation processes.

Solutions designed to meet these challenges have been available for more than a decade. But, although effective management systems have often been put in place for specific disciplines and IT resources, in most organizations enterprise-level integration remains a remote goal.

One reason for this is that solutions are themselves technically complex. Implementation challenges are magnified by the diversity of applications, platforms, and technologies in enterprise IT environments, and by the rate of change in these. Deployment initiatives must also deal with shifting business demands and with the need to create management processes that are highly interdependent and complex.

The most effective approaches have proved to be those that focus first on a subset of IT environments and then build toward broader structures. There are a number of potential candidates for this initial focus. But server bases are the obvious starting point.

In most organizations, servers account for a significantly higher percentage of IT spending than other infrastructure components. Even when they have been physically concentrated in data centers, server infrastructures remain dominated by small machines and diverse operating systems and software versions. Overall capacity utilization is typically the lowest, and downtime the highest for any major IT resource.

Early successes in server management thus offer the potential for larger savings and greater benefits in service quality than is the case for other segments of IT infrastructures. They may also act as a catalyst for, and may materially facilitate the realization of enterprise management goals.

The odds that this will occur are significantly increased if server management solutions are closely integrated and highly synergistic with those that address management challenges at the enterprise level. This is the case for Systems Director 6.1.

In addition to high levels of integration with IBM Tivoli solutions, which are discussed later in this report, Systems Director 6.1 interfaces to the industry's major third-party enterprise management offerings.

This approach, illustrated in figure 4, ensures functional consistency and technical compatibility between the enterprise and server levels of the management environment.

Figure 4
Systems Director Relationship to Enterprise Management Solutions

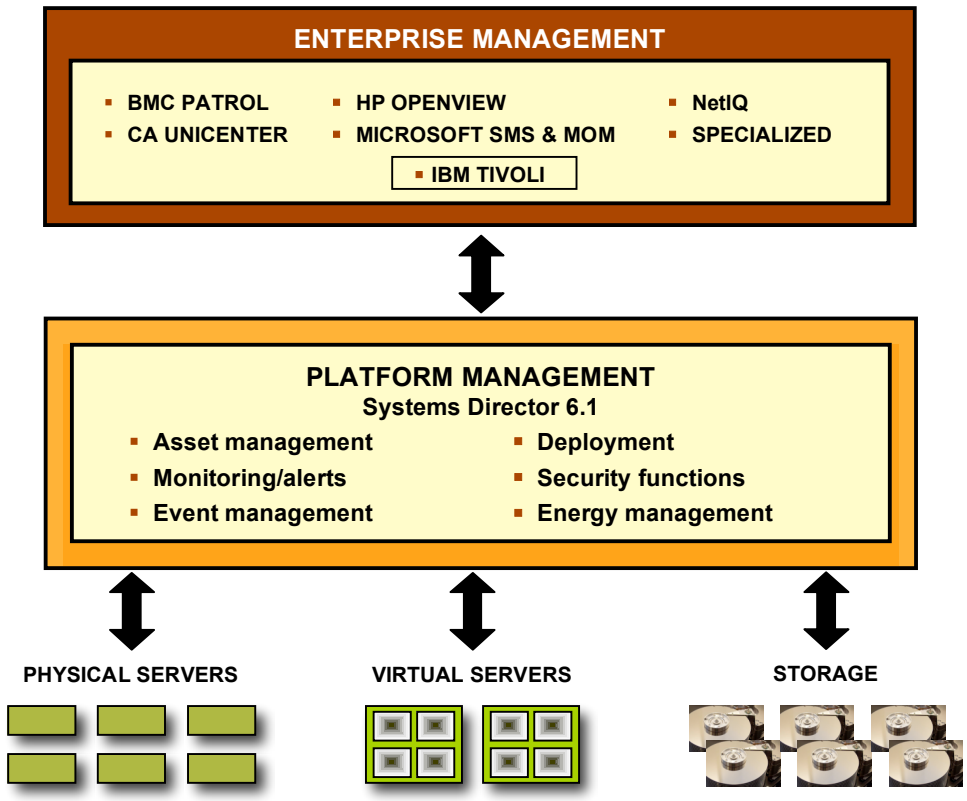

Third-party enterprise management offerings to which Systems Director may interface include BMC Software's Patrol, Computer Associates' Unicenter TNG, Hewlett-Packard's OpenView, Microsoft's Systems Management Server (SMS) and Microsoft Operations Manager (MOM), and NetIQ tools. Interfaces to specialized third-party tools, such as Altiris Deployment Solution, are also supported.

Conclusions: Saving the Lobster

There is an old saying that a lobster can be boiled without realizing it. As the temperature increases only one degree at a time, the lobster does not realize what is occurring until it is too late. Arguably, something similar has occurred in the IT world.

The growth of infrastructure inefficiency has been a gradual process. As organizations moved away from traditional centralized computing models in the 1980s and 1990s, IT infrastructures became dominated by small servers. The impact of the Internet and intranets and the growing complexity of software architectures and applications in the late 1990s and 2000s reinforced this trend.

More widespread use of management tools and consolidation initiatives has led to some improvements. In most organizations, however, physical server bases remain fragmented and are characterized by levels of inefficiency that would not be tolerated in other areas of the business.

Server virtualization offers part of the solution. But to pre-empt the effects of virtual server proliferation, virtualization must be coupled with aggressive consolidation and early deployment of advanced server management tools and processes.

System z, along with IBM System x and BladeCenter equipped with VMware and equivalents, and Power servers equipped with PowerVM offer significant advantages as consolidation platforms. Systems Director 6.1 is the obvious candidate to provide a common management solution for IBM servers.

If these are deployed at an early stage, major gains in infrastructure efficiency may be realized across physical as well as virtual resources. Organizations may also establish server environments that will ensure that future growth occurs in a cost-effective and non-disruptive manner.

The tools and technologies to realize such gains are becoming available. What is needed now is the ability to recognize and act upon the opportunities that they represent.

Additional Information

This ITG Executive Summary is based upon results and methodology contained in a Management Brief released by the International Technology Group. For copies of this Management Brief, please email requests to info-itg@pacbell.net.

International Technology Group

4546 El Camino Real, Suite 230
Los Altos, California 94022-1069
Telephone: (650) 949-8410
Facsimile: (650) 949-8415
Email: info-itg@pacbell.net