

## Mass Distributed Server Consolidation

### System z Mainframe Linux-on-z/VM Extreme Virtualization far Outclasses Over-hyped x86/x64 Approaches

#### About this Executive Summary

*This standalone Executive Summary abridges the main findings/conclusions from our recent, in-depth, 82-page White Paper of the same title, first published in February 2008.*

*Scale-out distributed servers (RISC-UNIX and x86/x64) vastly proliferated for a decade, with 28M in use by end-2007. These drove system management IT costs up six-fold (to a vast \$120B by 2007), and IT electrical power/cooling costs skywards (to \$22B by 2007). They have proved **economically catastrophic** and **environmentally abusive**.*

*Radical consolidation of this distributed computing scourge is now urgently required by 1000s of larger IT users. Once fashionable, such legacy, scale-out distributed server estates are now deeply discredited for their appalling resource waste and sky-high overall TCO.*

*Virtualization alone can eliminate this legacy distributed systems mess. The IBM System z mainframe remains unquestioned and absolute IT industry virtualization "Gold-Standard" and champion. One top System z (z9 or new z10), running Linux on the legendary z/VM extreme hypervisor, can now easily consolidate up to a thousand real distributed server (x86/x64-UNIX) workloads, bringing huge cost and resource savings and rapid ROI.*

*Our White Paper deeply assessed how and why such System z extreme virtualization & mass consolidation solutions are now far superior to palliative x86/x64 "fixes" (e.g. VMware ESX Server) currently proposed by "Wintel" advocates.*

*Detailed results, from 4 major IT users who are consolidating 4,700 distributed servers onto 37 System z mainframes, quantify and validate the impressive benefits now readily obtainable.*

#### Executive Summary

Our principal research findings/conclusions are highlighted in this Executive Summary, which was drawn/abridged from the extensive, in-depth analysis/discussion of our main White Paper of the same title.

**1. Distributed Computing = User Catastrophe:** Distributed servers became an **economic, environmental, and manageability catastrophe** for users! Server and software utilization averages 5-7%, one staffer is needed to support every 20-25 distributed servers, and many other severe issues, afflict these sprawls. Their **extreme inefficiency, gross waste of hardware/software capacity, excessive staffing, profligate electrical power/cooling costs, and extravagance with costly data center floorspace**, are now (*belatedly*) clear to all, more fully summarized in Figure 1 (*on page 2*). From the green environment/carbon footprint viewpoints, distributed systems were surely an unspeakable abuse. Strong words, but our Paper's depth analysis amply verified their truth.

**2. Software Strategies 2008 Systems Market Mega-trends:** With the credit crunch biting, markets falling, and economies now worsening, what are the top dozen early-2008 systems market mega-trend issues/directions? Figure 2 (*on page 3*) provides Software Strategies' ranking. IBM's System z mainframe, today's strongest solution to this "distributed server disease", centrally addresses, exploits, or delivers on, ten of twelve 2008 mega-trend issues, good cause for our Paper's System z focus.

---

*The most powerful "enabling technology" to optimize overly-complex, distributed IT infrastructures is virtualization.*

---

**3. Virtualization – Foundation Infrastructure Optimization Technology:** The most powerful "enabling technology" to optimize overly complex, distributed IT infrastructures is virtualization, which has recently (*and rightly*) grabbed IT industry and media center-stage. Although applicable across all types of IT-based system, we focus on commercial IT server virtualization.



## What's Wrong With Scale-out Distributed Computing Using Volume x86/x64 and UNIX Servers?




Figure 1: What's Wrong With Scale-out Distributed Computing Using Volume x86 (& UNIX) Servers?

- 4. Server Partitioning & Virtualization:** Sharing host server resources over multiple workloads/Operating System (OS) images/applications, each isolated in a partition, is the goal. Dynamic partitioning, where partition resources can be changed "on-the-fly", is vital where demands vary rapidly. Two hardware partitioning (*physical and software logical*), and two hypervisor-based (*Type 1: bare-metal and Type 2: hosted*) software models are seen today. "Type 1 bare-metal hypervisor" solutions (*used for IBM System z's dLPAR & z/VM extreme virtualization hypervisors*) are greatly superior but much the most difficult to engineer. Of the main ways to implement a hypervisor, "direct hardware support" (*used by System z*) is much the best, and "translate, trap & emulate" (*used by VMware ESX*), needed in x86/x64 virtualization, suffers considerable performance impact.
- 5. Virtualization/Server Consolidation Benefits:** Done well, large-scale, distributed server consolidation through virtualization delivers five hugely-valuable benefits:
- **Much-reduced management/support staff costs.**
  - **Much-improved flexibility and responsiveness.**
  - **Much-reduced software costs.**

- **Environment cost savings/green IT benefits.**
- **Sometimes reduced server hardware costs.**

We found these benefits are far strongest when the industry virtualization champion System z Linux-on-z/VM approach is used. They are much more limited when restrictive x86/x64 virtualization software-based approaches (*VMware ESX Server the generic*) are used.

- 6. IBM System z Mainframe – Virtualization "Gold-Standard":** IBM's System z is today's unquestioned virtualization "Gold-Standard", built via IBM's 40 years of mainframe virtualization pioneering/innovation. Select highlights include:
- **Dynamic LPAR Partitioning Hypervisor:** Rock-solid, <60 LPARs/system, deep z hardware support, and ultra-secure EAL 5 rated.
  - **z/VM Extreme Virtualization Hypervisor:** Hosts mid-100s to 1,000 significant-weight Linux virtual servers/System z.
  - **CPU, Memory, I/O, & Network Capacity:** Flexible, dynamic allocation of all z resource to LPARs & virtual servers.

## Systems Market – 12 Top Mega-trend Issues 2008 System z Mainframe Offers Top Solutions

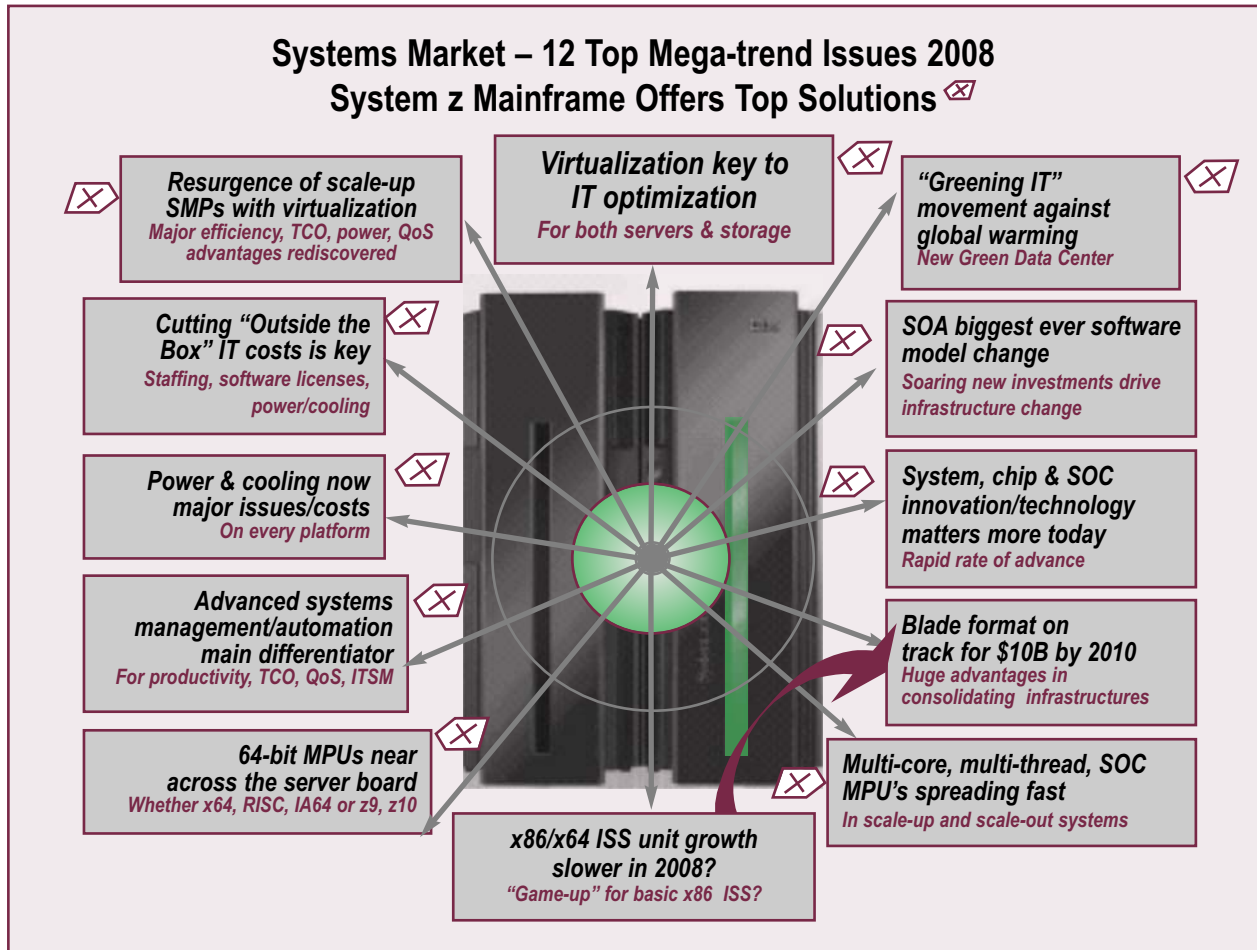


Figure 2: Systems Market – 12 Top Mega-trend Issues 2008 – System z Mainframe Offers Top Solutions

- **HiperSockets™**: High-performance, low-latency, internal System z TCP/IP network links between LPARs.
- **Intelligent Resource Director (IRD)**: SLA-policy-driven, cross-LPAR workload optimizer gives high z utilization and consistent SLA attainment.
- **Virtual Machine Resource Manager (VMRM)**: Manages resources applied to workloads over virtual server groups within z/VM.

*IBM’s premier “distributed server disaster/disease cure” is its System z mainframe, with the powerful Linux-on-z/VM extreme virtualization/consolidation solution.*

7. **System z Linux-on-z/VM Extreme Virtualization**: IBM’s premier “distributed server disaster/disease cure” is its System z mainframe, with its powerful Linux-on-z/VM extreme virtualization/consolidation solution. These together, IBM argues, deliver superior distributed server consolidation for many qualified workload types. Our White Paper found that the System z z/VM combination offers dramatic advantages when used to “sweep the floor clean” of such distributed server sprawl.

8. **Little VMware, Extraordinarily Large Hype**: First-to-market (2001) with a x86 server virtualization hypervisor, and later with related Virtual Machine Monitoring (VMM) management software, was specialist ISV VMware. Resold today by most x86/x64 vendors/channels, VMware became this segment’s “generic brand” for virtualization. Although sound and useful, this VMware ESX Server stack suffers major limitations/disadvantages (compared to the System z Linux-on-z/VM approach), as well as steep costs. It also now faces numerous direct competitors.

9. **Mainframe Transformation + New Workloads = Strong Resurgence**: The IBM mainframe underwent total transformation of its technology, economics, capacity/capability, openness, and in its software stack, from 1993 to date. Modern, new-to mainframe workloads, including Linux, Java EE™, SOA, ERP, distributed server consolidation, and BI/DW, were resurgence drivers, devouring 70% of post-2001 MIPS sold. Market revival began with 2000’s z Series, swelled to resurgence from 2004 (z990 & z890), and was amplified further by the System z9 generation mentioned below.

**10. 2005-08 System z9 Generation Highly Successful:** The System z9 generation successfully extended the mainframe's now strong resurgence for its long, two-and-one-half-year cycle, until the new System z10 generation arrived (see points 12 and 13 below). First seen as high-end System z9 109s (September 2005), the enhanced top-end System z EC and the lower-entry point mid-range System z9 BC soon followed (GA July 2006), with outstanding virtualization capabilities. Family highlights were:

- **System z9 Business Class (z9 BC):** For medium business. \$100K up, two models at 26-1,786 MIPS capacity.
- **System z9 Enterprise Class (z9 EC):** For large enterprises. 5 models, up to 54-Way CP, at <17,800 MIPS capacity.
- **New zIIP (System z9 Integrated Information Processor):** New specialty processor for low-cost DB2, network and XML workloads, on all z9s and z10s.
- **Unique "Hybrid Processor" Architecture:** Includes standard CPs, specialty Integrated Facility for Linux (IFL) (Linux), zIIPs (DB2, networks, XML), and zAAPs (Java and XML), mass I/O & channel MPUs, plus crypto MPUs, blended within these "commercial supercomputers", offering huge capacity in a modest footprint.
- **New FICON Express4 4Gbps Channel Speeds:** For improved 2X I/O capacity and performance on z9s.
- **Superb Virtualization/Automated Workload Management:** Enabling up to near-100% utilization of capacity.
- **Fullest Capacity on Demand (CoD) Options:** For maximum flexibility for customers to rapidly meet varying workload peaks.
- **Leanest, Greenest IT Platform:** Far the greenest solution, smallest data center footprint, lower power and cooling costs per 1,000 users, per virtual server, per workload, etc. (Versus distributed servers of equivalent capacity.)

**11. Lowest Cost Platform:** Flatly contradicting old myths, a decade of radical IBM mainframe hardware/software cost reductions, extreme-virtualization-driven efficiency, extensive automation/low staff needs, and excellent "green IT" strengths, now often make System z the lowest cost platform for substantial, mixed commercial workloads today.

**12. Spectacular New IBM z6 MPUs:** In late 2007, IBM previewed (to IT analysts) its spectacular, new **z6 System z CISC MPU**, that now powers its new, System z10 Enterprise Class (EC) next-generation mainframe, announced February 26<sup>th</sup>, 2008. These 4.4GHz. + ultra-high-frequency, quad-core z6 MPU chips pack 991M transistors, and use an "extreme-CISC" 894 instruction ISA (50 new in the z6).

Together with a huge, new 1.7B transistor, **System z SMP Hub** companion chips, these z6 MPUs have now further extended System z10 capabilities and capacity by wide margins. (See point 13 below.)

---

*...next-generation high-end System z10 EC mainframes that deliver up to 70% higher top-end capacity...*

---

**13. New IBM System z10 Mainframes Offer 70% More Capacity:** These highly-impressive z6 chips now power IBM's new, next-generation high-end System z10 EC mainframes that deliver up to 70% higher top-end capacity (new System z10 Enterprise Class (EC) – model S64), further extending System z's mass-consolidation capabilities. Lower z10 EC models span capacity points similar to their z9 EC forerunners but in denser form, using fewer cores and CEC parts. The still-competitive, mid-range System z9 BC continues as the entry-midrange mainframe for the present. (We expect a System z10 BC-z6 MPU-powered, replacement by early 2009.)

**14. System z – A "Data-center-in-a-box":** System power and extreme virtualization enable mass consolidation of (qualified) distributed server workloads onto this pre-integrated, pre-tested, complete System z "Data-center-in-a-box", shown in Figure 3 (on page 5), that offers:

- **Hundreds of Processors of Power:** Up to 416 total in a top-end System z9.
- **Extreme I/O Bandwidth:** And ultra-fast context switching.
- **Built-in Internal Networking:** High-performance, memory-speed, internal virtual networking, switching, & communications.
- **"Shared Everything" Resource Model:** Enables favorable Total Cost of Ownership (TCO), increasing returns to scale.
- **Product of \$B Engineering & Software Development:** Sustained, decade-long, massive \$Bs IBM R&D effort.
- **Highest Reliability, Availability, Security:** Of any commercial IT platform.
- **Lowest Staffing, Software, Space, Power, & Cooling Costs:** Combine to deliver lowest TCO when heavily loaded.

**15. IBM z/VM Extreme Virtualization Hypervisor:** A centerpiece of System z's "Gold-Standard" virtualization is its z/VM V5.3 extreme virtualization hypervisor, latest release of the industry's most sophisticated, scalable, and long-established, bare-metal hypervisor. Now ideal for mass (hundreds to a thousand) distributed server workload consolidation onto Linux-on-z/VM, the hypervisor has long also been used by IT groups to host/test multiple OSs, releases, etc., on a single mainframe.

## System z “Datacenter in a Box” The Escape From Distributed Systems Nightmares



### System z Mainframe Offers:

- Hundreds of processors
- Extremely large I/O bandwidth
- Built-in internal networking
- “Shared everything” resource model
- \$B engineering and software development
- A pre-integrated data center in one box
- Highest availability, security
- Lowest total, staffing, software, space, power, and cooling cost

*Using extreme virtualization with Linux-on-z/VM*

### Distributed Systems Suffer From:

- Highest costs, and highest complexity
- Masses of servers, storage, network gear
- Low hardware/software utilization
- Data silos and synchronization issues
- Highest, linear staffing costs
- Frequent outages
- Severe security issues and, costs
- Highest space, power and cooling costs




Figure 3: System z “Datacenter in a Box” – The Escape From Distributed Systems Nightmares

*A centerpiece of System z’s “Gold-Standard” virtualization is its z/VM V5.3 extreme virtualization hypervisor...*

**16. IBM Systems Director Vital Part of Solution:** IBM Systems Director, now a highly-capable, elegant, and comprehensive systems manager for physical and virtual platforms, manages servers, storage, and networking equipment, across all IBM server and storage systems, and major third-party platforms. It complements z/VM on System z with its z/VM Center component.

**17. Freedom From Distributed Computing Scourge With System z:** About 28M distributed x86/x64-UNIX servers are running today, some 50% (14M) on workloads better consolidated onto System z mainframes. Not many more than 14,000 large IBM z10 mainframes could, we estimate, take over/consolidate all these 14M distributed server workloads (@ <1,000 virtual Linux servers/System z10). This would rescue users from half of today’s distributed computing scourge, as well as saving these users scores of \$Bs in costs over the next decade.

**18. System z IFL Specialty Processor for Linux Enabler:** Underpinning soaring Linux on System z workload growth were IBM’s IFL specialty processors (2000 debut). These run Linux workloads only, have no software-cost impact, and are priced far below standard z9 or z10 CPs (@ \$95K for z9 BC, & @ \$125K for z9 EC and z10 EC).

**19. System z All-IFL, Linux-only Mainframes Enabler:** All-IFL Linux-only System z mainframes (also 2000 debut) leverage IBM’s highly-favorable IFL pricing, and eliminate regular System z software charging, to provide a highly-competitive Linux HW/SW platform. Ideal for mass distributed server to Linux-on-z/VM extreme consolidation, such machines could already scale-up to the 17,800 MIPS capacity of top regular z9s (System z9 EC, S54 54-way). Each of these could already easily support <400-500 mid-weight Linux virtual servers. A realistic example of such z9 consolidation is depicted in Figure 4 (on page 6). Now, with the new System z10 top-end S64, 64-way machine offering 70% greater capacity, an all-IFL z10 S64 can consolidate 700-850 similar-weight Linux VS images.

## Mass Consolidation of Distributed Scale-out Servers onto System z Mainframe Linux-on-z/VM – Example

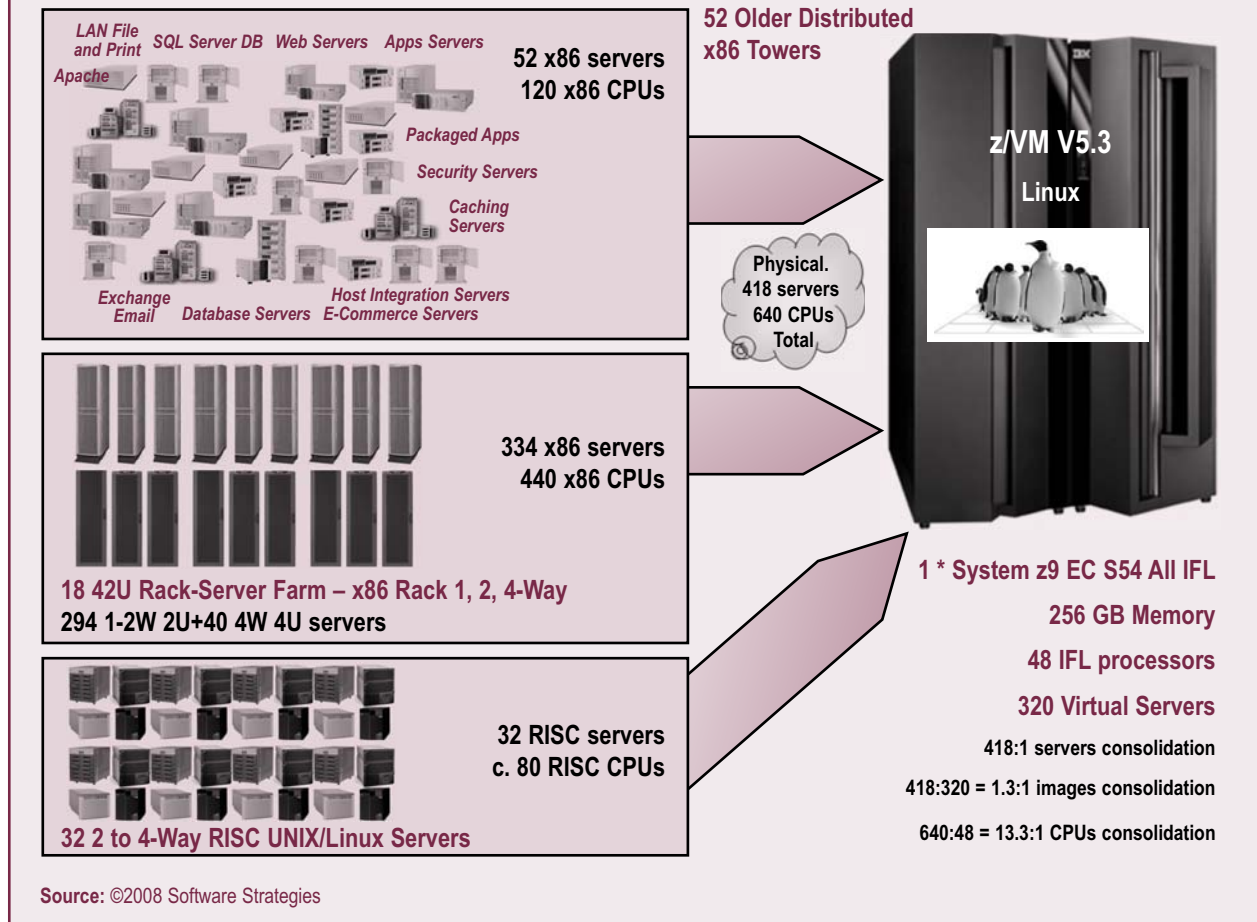


Figure 4: Mass-consolidation of Distributed Servers onto System z Linux-on-z/VM – An Example

### 20. System z Linux – Success far Above Expectations:

Since 2000, Linux on System z far exceeded expectations, with 2,500+ mainframes already running the OS (16%+). Usually run in a LPAR or under z/VM, Linux has today become an enterprise-grade OS, complementing System z's unique hardware Quality of Service (QoS). The OS brought many hundreds of new, originally UNIX, Linux, or Windows, software applications, middleware and tools, software products onto System z. When combined with z/VM (point 15 above), IFLs (point 18 above), and All-IFL Linux mainframe hardware systems (point 19 above), Linux provides the perfect OS foundation for mass distributed server consolidation.

### 21. z/VM – Leadership Features and Capabilities:

Our White Paper evaluated and assessed the features and capabilities of z/VM and its rich supporting IBM software infrastructure at length.

We found this outstanding, extreme hypervisor offers a tremendous depth and breadth of capabilities, far beyond others, built and refined from its predecessor's decades of heavy-duty usage. Our Figure 5 (on page 7) highlights just what we found to be the principal leadership features offered by the latest z/VM V 5.3 release.

### 22. Strong IBM z/VM Support Software Portfolio:

Complementing the extensive, inclusive capabilities of z/VM V5.3, summarized in Figure 5 above, our assessment also found that IBM now offers a strong portfolio of complementary, integrated software tools for the platform. These include:

- **IBM System Director:** Physical & virtual server, storage, networks manager for all-IBM and other platforms.
- **IBM TotalStorage Productivity Center:** IBM's extensive physical/virtual storage management suite.

- **IBM Tivoli Provisioning Manager:** Provisions OS & middleware software images onto z Linux VS under z/VM.
- **IBM Tivoli OMEGAMON XE for z/VM, Linux:** Unified z/VM & Linux System z performance monitor.
- **RACF Security Server for z/VM:** Improves z/VM V5.3 data security, & controls user-allowed actions.
- **IBM Backup-Restore Manager for z/VM V1.2:** Backs up & restores z/VM & Linux system files & data.
- **IBM Tape Manager for z/VM V1.2:** Manages, monitors, & automates tape resources/operations.
- **IBM Operations Manager for z/VM V1.2:** VM monitor/manager that automates maintenance tasks & responses.
- **IBM Archive Manager for z/VM V1.1:** Helps users archive historical/infrequently used data.
- **GDPS/PPRC Multiplatform Resiliency for System z:** GDSP business resiliency for z/OS & Linux on z/VM.

**23. IBM z/VM V5.3 vs. VMware ESX Server V3.1:** Our full White Paper published detailed comparisons, deeply examining both these virtualization platforms from our distributed server consolidation viewpoint. One analogy neatly summarizes our comparison results. The Linux-on-z/VM System z solution compares to the VMware ESX Server solutions on x86/x64 like a **Boeing 747 Jumbo jet compares to a fleet of 100 4-passenger Cessna twins** as a means of air passenger transport. Both means, in theory, could be used to transport many people long distances. Where the mission is to carry 400 passengers, at the same time, on a 6,000-mile leg over water, economically, rapidly, and safely, the 747 wins easily. It alone had the capacity (*400 seats*), range (*8,000+ miles*), and far the lowest staffing (*16 vs. at least 100*).

It also pays only one set of take-off/landing fees (*1 vs. 100*), can fly in most weathers, has an outstanding safety/security record on such flights, and delivers the lowest mission fuel cost, and total cost, per seat-mile.



Figure 5: z/VM Extreme Virtualization – Leadership Features

z/VM on System z (*the 747 in our analogy*) is the heaviest-duty, most industrial-strength, highest-productivity, lowest-unit-cost, longest-perfected, and most trusted extreme virtualization hypervisor in the IT industry, and the only one today capable of mass distributed server consolidation! The Cessna machines (=VMware ESX Server) are excellent light planes, designed/optimized for shorter leisure or business flights, usually within one nation, and of a few hundred miles length, for 2-4 passengers only, in fair weather.

**24. 4,700 Distributed Servers (DS) to 37 System z Mainframes:** 4,700 is the industrial-scale number of distributed servers consolidated onto 37 new System z mainframes, at four actual mainframe user organizations listed below:

- # 1. IBM Project Big Green – 3,900 distributed servers onto 33 System z, in work.
- # 2. Nationwide – 500 distributed servers onto 2 System z, to date.
- # 3. Government of Quebec – 200 distributed servers onto 1 System z, so far.
- # 4. Nexsar Group – 80 x86 distributed servers onto 1 System z9 BC, so far.

These diverse, major examples, fully assessed in our White Paper, validated the case and qualified the impressive benefits in 25 below.

**25. Dramatic Benefits for System z Consolidation Users:** Impressively large and consistent System z Linux-on-z/VM benefits were delivered, according to these users' distributed server consolidation experiences, and other data sources. These benefits include:

- Reducing server numbers by 100-200: 1 System z, or more.
- Raising CPU/hardware utilization, 10-20:1 times higher, up to 85-95% on System z from 5% on x86/x64.
- Supporting far more virtual servers/CPU cores, by 15-30:1 System z IFL core.
- Far fewer CPU cores needed, by 10-20:1 System z IFL core factor.
- 10-20:1-fold fewer CPUs worth of SW licenses to buy, 4-8:1 lower software cost.
- 4-8:1 fewer support staff needed, and the same ratios of lower support staff cost.
- 1 support FTE/100-200 System z Linux virtual servers, vs. 20-25 physical distributed servers/FTE.
- Hardware costs actually now often comparable.

- Dramatic power & cooling cost savings, 10-20:1-fold less cost, often 90%+ reduction.
- Data-centre floorspace, often 25:1-fold less – e.g. slashed from 10,000-sq. ft. (DS) to 400-sq. ft. (System z).
- Networking equipment/cabling cost savings, 3-4:1 average reduction.
- Far faster server provisioning, 15-1,000:1-times faster; seconds/minutes and not days/weeks.
- Enables technically-feasible, affordable Disaster Recovery/Business Continuity (DR/BC), impractical with distributed servers.
- 30-50% overall TCO savings consistently reported.
- Rapid ROIs with paybacks in the 15-months-to-2.5-years range.

These points highlight just how much wastage distributed systems incur, and just how dramatic the major benefits and saving of System z Linux-on-z/VM mass consolidation are today.

---

*Such huge savings cannot be ignored, painting a compelling business case for System z.*

---

**26. Huge \$TCO Savings Available – \$114M and 51.9% Saving Example:** Another real-customer-based, detailed 5-year TCO comparison was included in our main White Paper (*Figure 21 Page 57 there*). This evaluated replacing 2,239 existing suitable distributed servers by five new all-IFL System z9 mainframes running z/VM, which showed \$114M/51.9% 5-year TCO savings over their existing distributed solution. Such huge savings cannot be ignored, painting a compelling business case for System z.

## About Software Strategies

Software Strategies is a specialist analyst firm focused on enterprise IT platform strategies and issues. Specialist expertise on mainframes, servers, operating systems, and on middleware software/tools, has been our common thread. Since 1997, we have worked closely with numerous industry leaders, including: IBM; Unisys; Microsoft; Intel; Misys; Fidelity National Information Systems; CA; BMC; Stratus Computers; ICL; NetIQ; and others. Many tens of thousands of Enterprise IT users have benefited from our authoritative reports, white papers, and our presentations at scores of IT events, seminars, and conferences.

## About the Author



**Ian Bramley**  
Managing Director – Software Strategies  
[ian.bramley@software-strategies.net](mailto:ian.bramley@software-strategies.net)

This standalone Executive Summary was researched/written by Ian Bramley, Managing Director of Software Strategies, and was published in March 2008. (It summarizes the firm's major, 82-page White Paper under the same title, the 2nd Edition of which was published in February 2008). The views expressed are those of Software Strategies alone, and are based on our proprietary research. Ian founded Software Strategies in 1997. He is an experienced enterprise infrastructure analyst, has published scores of popular reports and white papers, and has served as a keynote speaker at many industry events. Previously, he was Director of Enterprise Platforms at Butler Group, and Founder/Chairman of the Enterprise NT Management Forum industry group from 1998 to 2001. Previously, he held executive positions with four international software/services vendors over a 25-year prior IT industry career.

**1<sup>st</sup> Edition Published:** March 2008

Design & Layout: iok design – [info@iok-design.co.uk](mailto:info@iok-design.co.uk)

**Important Notice:**

The information available in this publication is given in good faith and is believed to be reliable. Software Strategies expressly excludes any representation or warranty (*express or implied*) about the suitability of materials in this publication for your purposes and excludes to the fullest extent possible any liability in contract, tort or howsoever, for implementation of, or reliance upon, the information contained in this publication. All rights reserved. This publication, or any part of it, may not be reproduced or adapted by any method whatsoever, without prior written Software Strategies consent. Software Strategies acknowledges all trademarks referenced in this paper.