Arcati Mainframe Yearbook 2024

Arcati Mainframe Strategy Papers

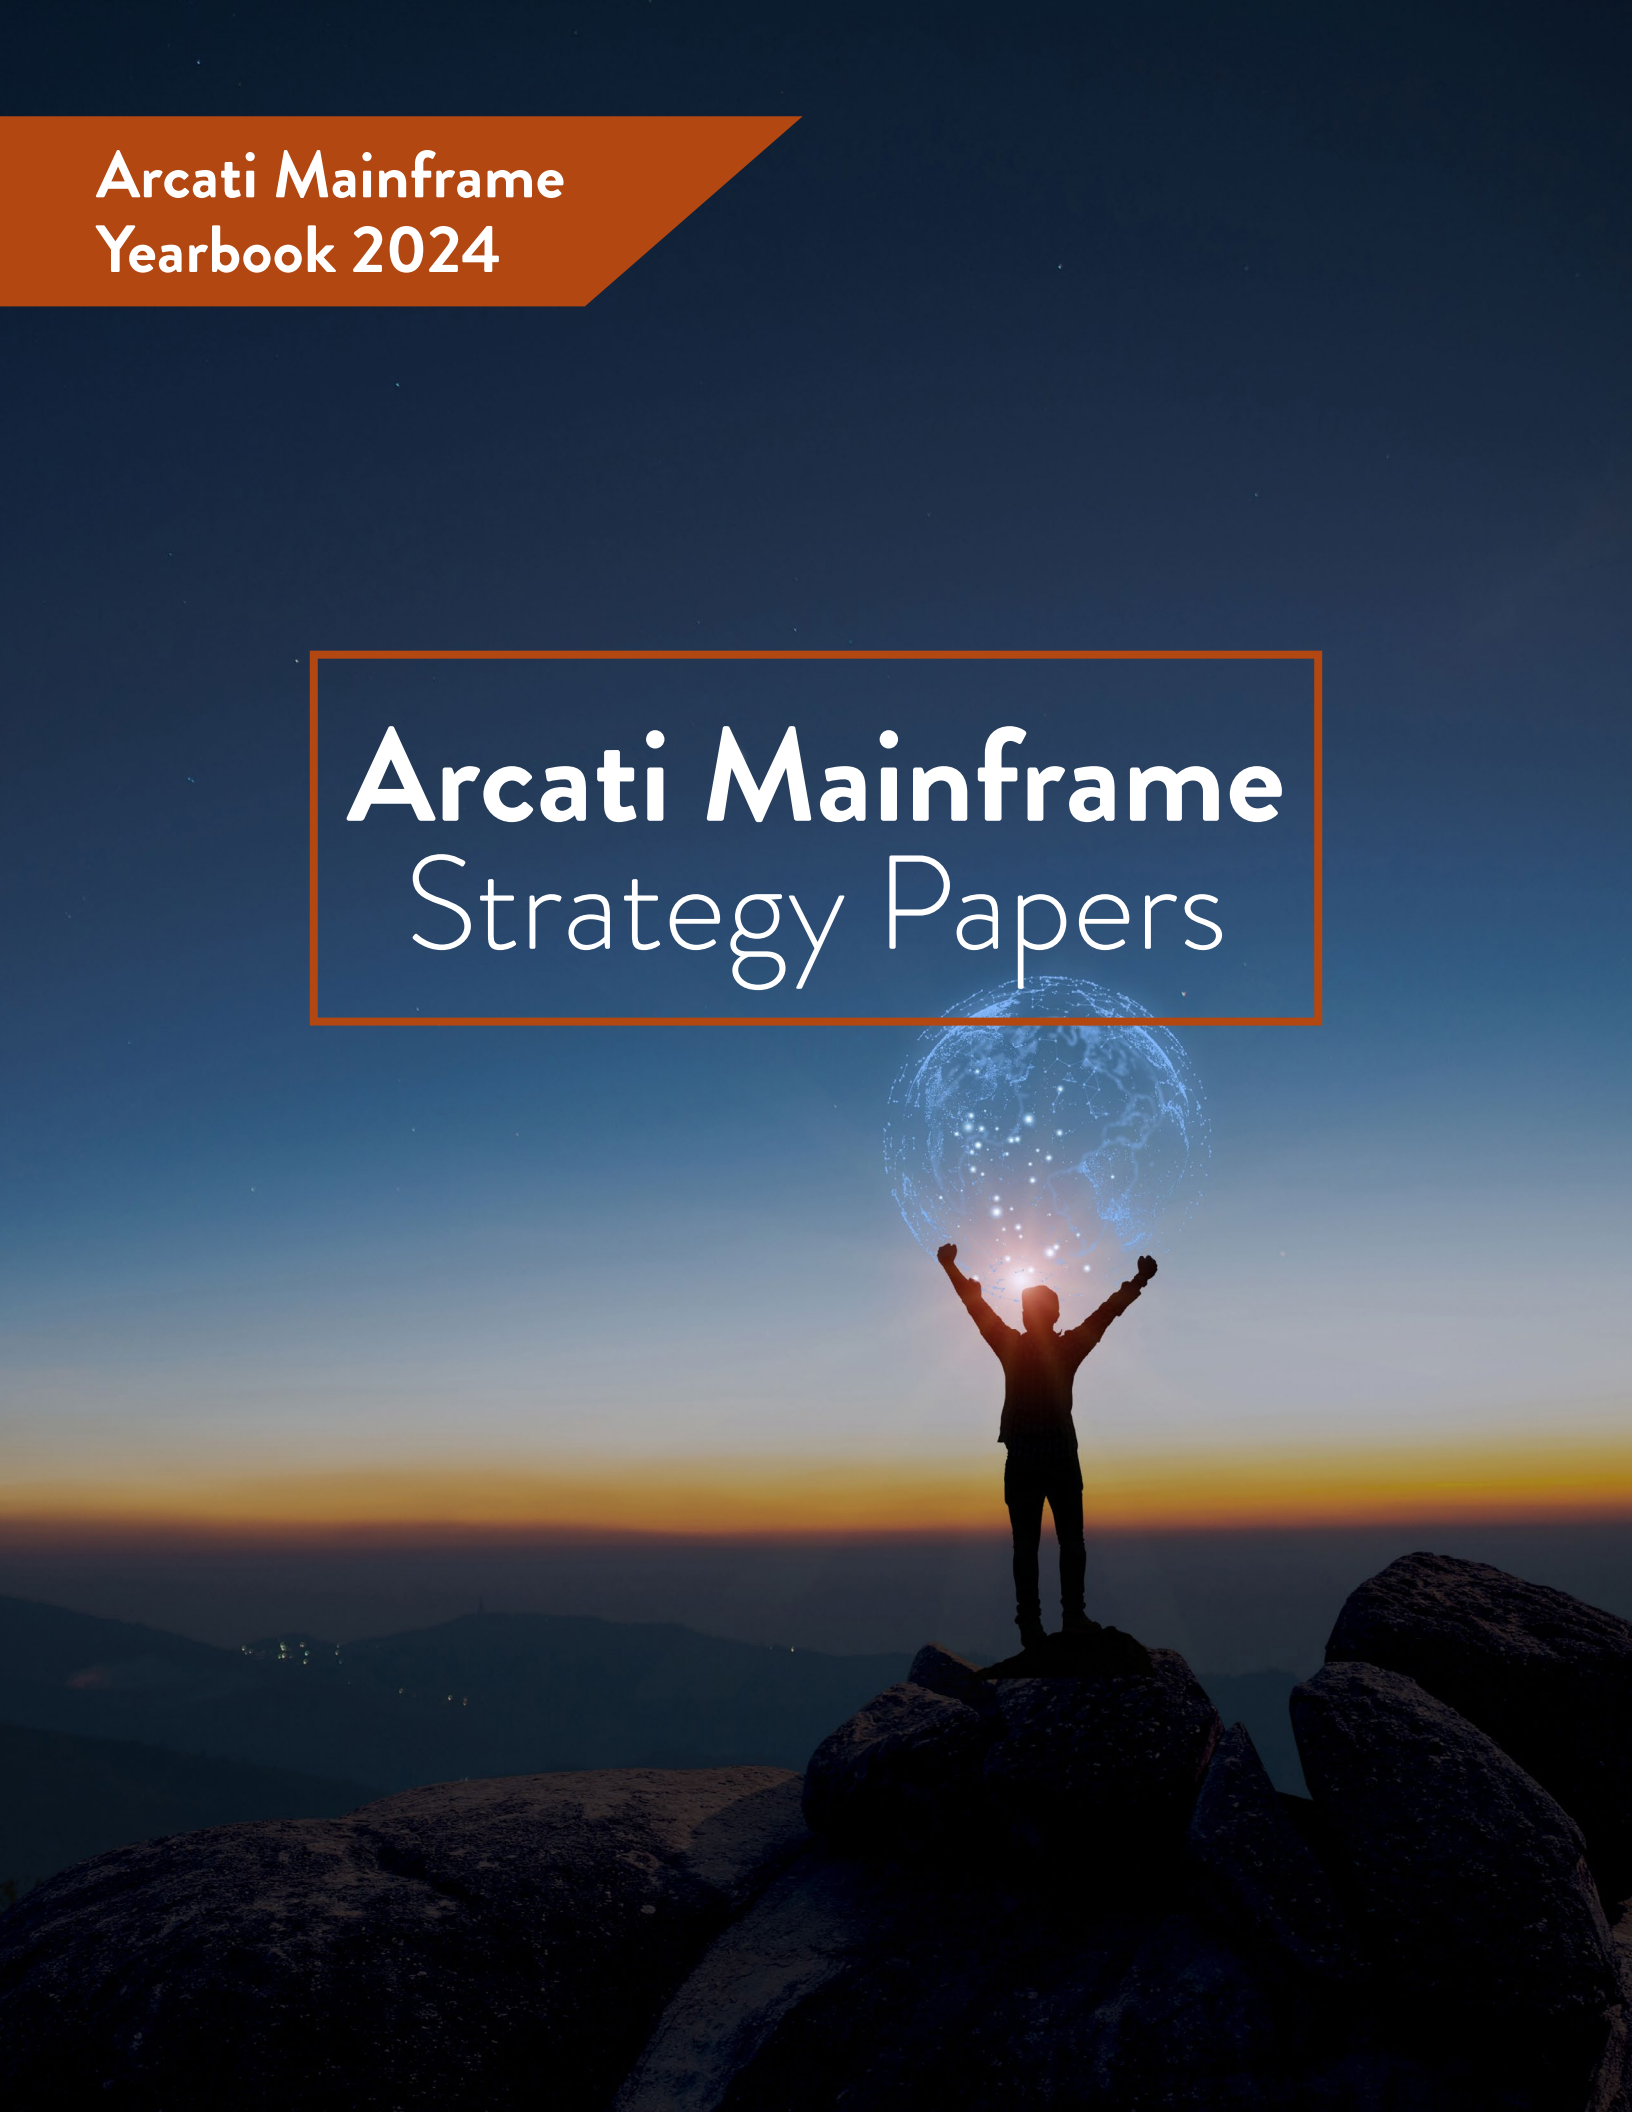

Mainframe is a Part of Your Cloud Strategy.

Now What?

Three Ways to Include Mainframe Workloads in Your Hybrid Cloud

Matt Hogstrom

Distinguished Engineer, AIOps Automation & Cloud Integration, Mainframe Software Division

Broadcom Inc

If Mainframe were an athlete, it would have multiple MVP titles. So would Cloud. Now, this sounds like the beginning of a championship team. But, just as great players alone don't bring home the pennant (or trophy, depending on your sport), neither do great technologies. You've got to bring those "players" on your IT roster together and integrate them into a cohesive unit that allows each one to shine. That's what makes a championship team.

Now, in sports, players get an off-season and time to rest. In the world of business—and the technology that runs it—there is no off-season. Your IT stack needs to be ready 24x7x365. That means every one of the platforms on your IT roster needs to know its position and be able to perform at its best and work seamlessly with the rest of the squad.

Many companies want the flexibility that Cloud offers, yet business-critical workloads that depend on the Mainframe aren't going anywhere. In fact, there's great value in what the Mainframe delivers for business and operations. Activating the Mainframe as part of your Cloud strategy means opening up access to a rich data source—a wellspring of modern insights and applications that transforms your ability to build business innovations and IT resiliency. So the question is not so much should you include the Mainframe in your Cloud strategy, but how?

Go Hybrid, Go Mainframe

When creating a successful hybrid environment with the Cloud and Mainframe, it's important to understand that you are modernizing the infrastructure, not individual applications. A Cloud strategy alone only modernizes infrastructure for Cloud-like workloads. The Hybrid approach takes capabilities on the Mainframe that have proven value to your business and delivers them to all applications on the Cloud.

In general, a Hybrid Cloud architecture enables teams to:

- Consolidate and share IT resources
- Orchestrate processes with the help of automation
- Connect multiple systems through a network
- Scale and quickly provision new resources
- Incorporate a single, unified management layer
- Move workloads between environments

The Mainframe is the fastest and most secure platform on the planet and constantly expands based on evolving technologies and business needs. Combining the strengths of Mainframe with Hybrid architectures allows organizations to continue to leverage proven value and critical capabilities. The challenge is to select the best approach for your business and current IT stack.

Start by evaluating your current tech landscape against your business needs. This knowledge will enable you to identify value and differentiate which applications run best on the Cloud and which are better suited for the Mainframe. It's a winning lineup that helps you capitalize on the strengths of both.

Here are three proven ways to successfully integrate and benefit from Mainframe workloads in your Hybrid Cloud.

#1

Enable Cloud Access to Mainframe Data

The Mainframe hosts a trove of critical business records and data—offering impactful insights on everything from operations to customer experience. This data is a source of incredible value that Cloud apps can, and should, leverage to gain advantage in the market.

Traditionally, it could be challenging to access Mainframe data from outside the host. Modern applications such as online shopping and banking are primarily API-based and Mainframe is not. These incompatible formats can result in application projects that require Mainframe data taking a long time.

Today, businesses are using APIs to access Mainframe data with very positive results. APIs enable secure and managed access to the Mainframe and help abstract incompatible formats so that Cloud-native applications, such as a mobile banking app, can easily leverage valuable Mainframe data. Opening up the Mainframe with APIs means that businesses can combine the power of the Cloud and Mainframe to develop modern solutions with more agility and faster time to market. In addition, standard data integration technologies such as RESTful APIs, virtualization, or GraphQL make it easy for businesses to support flexible development.

#2 Enable Cloud Access to Mainframe Services

Using RESTful APIs, businesses can make accessing Mainframe services and capabilities look the same as they would on any Cloud service. This familiarity is quite handy. For example, Mainframe services usually come with business logic, policies, or processes around using or updating the associated data. These are part of standard business services required to comply with regulations or compliance criteria. Incorporating these services into applications from outside the host can be a challenge. However, now you can leverage the existing policies and logic with a RESTful interface.

A RESTful interface, synonymous with RESTful APIs, is how businesses operate digitally and manage interoperability between services and developers. In essence, RESTful APIs make accessing Mainframe services and capabilities operate just like they would for any Cloud service. Enabling Cloud access to Mainframe services in this way means you can more easily modernize and accelerate the delivery of Mainframe apps. For example, creating new customer-facing interaction logic to function in applications in real-time.

#3 Host Cloud Workloads on the Mainframe

The promise of Hybrid is mainly in portability and optimization—provisioning workloads where they make the most sense and moving workloads between platforms as needed. Modernizing your Mainframe infrastructure to run newer, Cloud-native workloads means gaining the ability to make those new languages and runtimes available on the Mainframe for developers to exploit.

Many newer workloads are non-traditional, as in not COBOL or PLI running in CICS or IMS. They include new databases and runtimes found on Linux (or Linux on z), new languages like JavaScript, Ruby, or Python, and new technologies like containers and Kubernetes. Hosting new runtime technologies on the Mainframe, like containers in zCX or Linux on z, makes it easier for developers to stay current with workload runtimes in a shorter time.

Hybrid Cloud with the Mainframe FTW

Season after season, organizations make a substantial investment in services, business logic, governance, and compliance on the Mainframe. Leveraging that investment doesn't mean reengineering the existing assets. It means modernizing access to those resources using the language of the Cloud.

Given the Mainframe's "triple threat" of unique strengths—scalability, security, and reliability—businesses need to consider integrating it into the Cloud a necessity.

There are multiple ways to integrate the Mainframe based on your business' priorities. Assessing your business and operational needs will guide you toward the best one(s). Whichever way you choose, embracing and exploiting Mainframe strengths in your Cloud strategy will expand the value of your IT investment and set your business up to win with a game plan to innovate, grow, and offer extraordinary customer experiences.

Click Here to Watch the Video

The IBM Mainframe is Still the **Real Deal**

Keith Allingham

DataKinetics' CEO Allan Zander's article appeared on the Planet Mainframe blog in 2021, and it made some waves in the distributed systems world—see Hacker News discussion. The IT folks without much knowledge of IBM mainframe architecture refuse to believe the numbers in the article, while others lambaste the article for being old. Some of the comments are fair, but as AWS, Google, et al, provide more powerful cloud/distributed solutions, IBM does the same with its mainframe platform.

About 12 months ago, I had a conversation with an IBM tech in their computing costs group responsible for analysis and cost comparison between existing IBM Z installations against proposed cloud replacement proposals given to IBM customers by cloud business organizations (AWS, Google, Microsoft, etc.). He was interested in the Planet Mainframe article because it closely mirrored the results that he and his team were seeing in the 2020-2021 time frame, as they researched and compared an IBM Z installation vs an AWS proposal.

As you can see, this article cites data that is now more than 5 years old, but the comparison is still generally accepted as valid. Until a new article with new data is available, here is the Planet Mainframe article that started all the fuss:

The IBM Mainframe: The most powerful and cost-effective computing platform for business

Allan Zander, CEO, *DataKinetics*

Many of we mainframe pundits have written about the robustness, power, perseverance, capacity and more importantly, the cost-effectiveness of the mainframe (Allingham, Sun, Peleg), including myself. But what about showing the superiority of the mainframe using real numbers, comparing it to other platforms? That requires a lot more work. Schroder and Olders show us some real-world numbers, but how about showing the ugly details? That's even more work, and fortunately, a couple of people have done that as well.

Michael Benson's Enterprise Executive article in 2016 did that—since then, distributed servers have come a long way (AWS, Google and a host of other cloud service providers), but so has the mainframe. In 2015, the top-of-the-line mainframe was the z13, an outstanding business machine; since then, successive machines, z14 through z16 (and counting) outperform it considerably on many levels –speed, transaction throughput, security, flexibility, and more. A main argument then, as now, is cost; and that's a losing argument right from the get-go.

Comparing Platform Costs

"Other platforms are cheaper..." This is the basic claim for most people interested in dumping mainframe systems in favor of commodity servers. The argument is simple: "Google, Amazon and Microsoft don't use mainframe systems at their back end, so why should anyone?" Fair question, but let's look at the premise first—are server farms less costly than the mainframe? Recently, Michael Benson did a study for Enterprise Executive magazine in an article called CIOs: Are You Really Paying Less by Using x86 Platforms? In it, he configured two similar performing platforms—one mainframe-based, using an IBM z13 mainframe system, and the other, a bank of HP servers. Table 1 shows the system specifications.

ATTRIBUTE	HP PROLIANT BL460 GEN9	IBM Z13 2964 N30
Total Servers	12	1
Processors	24	30
Cores/processor	12	1
Cores/server	24	30
Total cores	288	30
# VMs	1000	1000
Memory	2 TB	2 TB
Hypervisor	VMware vSphere 4	IBM z/VM
Cloud Mgmt	VMware vRealize	IBM Wave
OS	Red Hat Enterprise Linux	Red Hat Enterprise Linux
Web server	Apache HTTP	Apache HTTP
Application server	IBM Websphere	IBM Websphere
Messaging	MQ	MQ
Database	Oracle EE	Oracle EE

He explains that running Linux on the mainframe is no different than running it on x86 servers. The only real difference is the cost, and the belief is that x86 platforms do it for less. But do they? The hardware costs for these configurations run in at $2,299,451.00 for the server farm solution, and $2,793,371.00 for the mainframe solution. However, due to licensing costs, the software cost for the server farms comes in at $1,807,406.00, with the mainframe running at only $416,883.00.

ANNUAL LABOR	HP PROLIANT BL460 GEN9 (QUANTITY=12)	IBM Z13 2964 N30 (QUANTITY=1)
Server admin	$580,160	$430,000
Net admin	$384,000	$76,800
Total	**$964,160**	**$506,800**

So yes, the hardware is cheaper, but there is not quite as much difference as you might expect. The real surprise is the difference in software cost. When you also consider maintenance costs, the pattern continues. Maintenance

costs for the server farm come in at $390,327.00, with the mainframe at $269,767.00. Labor costs are also part of the picture.

At the end of the day, what really matters is the ongoing operational costs of the two platform solutions. Table 3 shows a considerable gap in favor of mainframe computing.

Over a five year period, operating costs compound, and the picture looks much worse for the server farm, $9,052,749.00 vs $6,979,693.00 in for the mainframe setup. The shocking conclusion therefore, is that it is cheaper to run the mainframe system than it is to run the server farm.

OPEX	HP PROLIANT BL460 GEN9	IBM Z13 2964 N30
Hardware mtce	$9,544	n/a
Software mtce	$390,327	$269,767
Admin	$964,160	$505,800
Other (power, etc.)	$31,505	$68,355
Total	**$1,395,536**	**$844,922**

When doing cost comparisons, it is good practice to look at all contributing costs, and to look at long term cost of ownership. This comparison would have looked a lot different if we stuck to just the hardware acquisition cost, or even if we hid the personnel costs in a general employee pool rather than in the TCO calculations.

Technology Economics

Cost is one thing—often a very misunderstood thing, as Michael Benson pointed out. But acquisition and ongoing cost represent only one dimension in a complicated cost-comparison between platforms. What about environments that run a mix of mainframe and distributed systems? And what about comparing not just cost between platforms, but real costs in specific industries? Well, that's where Dr. Howard A. Rubin of *Rubin Worldwide*, a technology economics research firm, comes in.

In his paper, *The Surprising Technology Economics of Mainframe vs. Distributed Servers*, Dr. Rubin explains that understanding computing platforms and their economic relevance in the context of their contributions to business performance is critical. This context provides a transparency that goes far beyond the basic economics of the costs of hardware and software acquisition or a TCO calculation. This is especially important when we consider that technology costs are a rising part of ongoing business operations expense.

It is good practice to look at all contributing costs, and to look at long term cost of ownership.

IT costs vs business revenue and cost

Technology costs relative to business revenue and operating costs vary considerably from one industry vertical to another. For example, in banking and finance, IT expense represents about 6% of revenue and just over 7% of business operating expense; compared to the retail sector, where IT expense represents just under 1.5% of revenue and just over 1.5% of business operating expense.

Cost of platform choice

Businesses have choices on how they will handle their processing needs and this typically comes down to the mainframe and server farms. The cloud is part of the latter solution. The reality is that any business that runs mainframe systems also runs server farms, so it is fair to characterize them as running "mainframe-heavy" datacenters, while those without mainframe run "server-heavy" datacenters. It is also useful to consider new metrics for these datacenters—MIPS per $1M of revenue and physical servers per $1M of revenue. These aren't equivalent in any way, but they serve to represent the economics of their computing choices in measurable economic terms.

When comparing businesses within the same industry vertical, the "heaviness" of their IT deployment strategies result in a significant differences. For example, for financial services businesses:

Mainframe-heavy shops consume:

- 3.1 MIPS per $1M of revenue
- 0.22 servers per $1M of revenue

While the server-heavy shops consume:

- 1.75 MIPS per $1M of revenue
- 1.2 servers per $1M of revenue

BUSINESS	AVERAGE COST	MAINFRAME-HEAVY COST	SERVER-HEAVY COST
Distribution	$4,255,273	$3,936,728	$6,809,818
Communications	$4,979,371	$4,306,000	$8,295,000
General business	$4,832,000	$4,414,000	$7,846,000
Computer Services	$6.093,958	$5,644,350	$7,619,000
Industrial	$9,270,513	$9,082,000	$11,181,000
Financial Services	$12,627,002	$12,742,000	$16,445,000
Government	$15,161,129	$14,148,000	$15,981,703
Average	**$8,174,178**	**$7,753,297**	**$10,596,789**

When these figures are mapped to the total cost of mainframe and server farm costs within various industry verticals, the economic differences that can be attributed to their deployment strategies become apparent—(Table 4). The inescapable conclusion is that mainframe-heavy computational costs to support a $1B organization on average may be 30% lower than a server-heavy deployment.

Cost of Goods

While the cost of technology yields interesting conclusions, the actual costs of platform choice are also surprising, and support the former. The next step is to link the technology costs to business costs.

A good way to do that is to use a cost-of-goods metric. Ask the question: "What is the IT cost contribution to the business cost of goods?" And follow that up with: "How does technology deployment affect the measure of impact on the business?" Table 5 itemizes the cost of goods for five business types—finance, industrial, communications, general business and insurance.

This data implies that where appropriate, effective use of mainframe resources results in a 29% cost advantage over distributed server-heavy deployments.

Looking closely at the insurance data, we see that the average IT cost of processing an insurance claim in a mainframe-heavy environment is approximately $56, which is $36 less than the processing cost in a server-heavy environment. What does that mean to an insurance business? For an insurer that processes 100,000 claims per year, the savings could be $3.6 million per year by leveraging mainframe technology.

Similarly, a bank with 4500 ATMs would be paying over $1000 per ATM using a server-heavy datacenter, as compared to less than $600 using the mainframe-heavy scenario. Such a bank could save more than $2 million per year by leveraging mainframe technology. To be fair of course, ATM costs are only one small part of a bank's IT cost concerns.

Competitive advantage

Any large company interested in maximizing computing power AND controlling costs will clearly enjoy a competitive advantage over a similar company that just seeks to avoid mainframe technology in favor of server farms. This advantage translates directly to the bottom line, shareholders and investors. And for a company considering a mainframe migration project as a means for cutting costs, this information could be seen as "found money."

PROCESSING COST PER:	AVERAGE COST	MAINFRAME-HEAVY COST	SERVER-HEAVY COST	RATIO, MF VS SERVER
ATM	$928.00	$572.00	$1,021.00	56%
SKU	$227.27	$184.09	$252.27	73%
Mobile subscriber	$23.26	$18.26	$26.12	70%
Patent	$390.83	$372.00	$401.00	93%
Claim	$78.00	$56.00	$92.00	61%
			Average	**71%**

Conclusions

The facts support the notion that the mainframe is the most powerful and cost-effective computing platform for large businesses with a need for high-intensity transaction processing. Claims to the contrary are typically either as a result of simple lack of knowledge on the subject, or a biased unwillingness to look objectively at the facts.

But if the mainframe is so great, then why is it not being used by the newest and latest concerns (Amazon, eBay, etc.)? The reason is bias. Whether intentional or through ignorance, there is a great deal of bias against the mainframe. We hear it all the time – and saw it in the comments to the original publication of the article. People say "It's too expensive!" (It clearly is not.) "It's old and dusty!" (Obviously not.) "It's hopelessly outdated!" (Not even close.) "I don't know very much about it!" (Ahhhh . Now we're getting somewhere.)

The last part is the key to the puzzle of why the mainframe generally has a difficult time displacing server farms in environments where it could make a positive impact. The truth is, organizations that could benefit from the mainframe, but don't, are leaving money on the table.

So, What's the Strategy?

We've told you what IBM already knows, what many IBM customers already know, what some "mainframe replacement" vendors secretly know, and even what today's big cloud vendors know. So, what's the strategy moving forward?

What about mainframe shops having trouble keeping up with growing workloads on their "most powerful and cost-effective" mainframes? Should they be upgrading? Shifting workloads off-platform? As you might guess, there are options. There are a couple of organizations that are helping mainframe shops to optimize what they have now—to increase workload throughput of the systems they're currently running. No upgrade needed; no changes to application logic, no changes to the Db/2 (or whatever) database being used. This is possible using high-performance in-memory technology.

And both IBM and DataKinetics are offering these solutions right now. Talk to people who actually know something about the platforms under evaluation.

DATAKINETICS
Z PERFORMANCE & OPTIMIZATION

Click Here to Watch the Video

Micro-Segmentation Keeps Sensitive Mainframe Data in Compliance

Provided by: *Vertali*

Executive Summary

Mainframes hold an organization's most critical and sensitive business data, making it crucial to ensure that data is secure and meets the strictest privacy regulations.

Controlling access through network micro-segmentation is an effective way to protect sensitive data on mainframes by isolating applications or devices. Such isolation is required in heavily regulated industries with compliance standards such as the Payment Card Industry Data Security Standard (PCI DSS), Health Insurance Portability and Accountability Act (HIPAA), and General Data Protection Regulation (GDPR).

Micro-segmentation is an important step toward achieving Zero Trust security. Micro-segmentation can isolate each application into its own network segment. That gives organizations the ability to limit application access to specific network segments or specific devices, providing an additional layer of security beyond user authentication.

Isolating card payment processing applications to specific network segments can greatly reduce the scope, cost, and time of PCI DSS compliance assessments. Although segmenting the cardholder data environment (CDE) from the rest of an organization's network is not a PCI DSS requirement, it is highly recommended by the PCI Security Standards Council. By consolidating data into fewer locations that have more control over that data, segmentation reduces the risk to an organization's payment account data.

The PCI Security Standards Council says that any assets that store, process, or transmit payment card data are "in scope"—meaning they must be assessed for PCI compliance. Thus, the entire network is in scope without proper segmentation. The wider the scope, the longer and more costly the PCI compliance problem becomes.

Network segmentation that isolates the card handling applications reduces the PCI review to that specific area rather than an entire network, which can span hundreds of thousands of devices. Reducing the scope of the PCI DSS assessment also reduces the cost and difficulty of implementing PCI DSS controls. It also mitigates risk to an organization by consolidating cardholder data into fewer locations with greater control.

Benefits and Challenges of Micro-segmentation

Segmentation divides a network into segments to make them easier to secure and manage. Micro-segmentation goes beyond that, carving out a segment for each application, isolating and containing the traffic within that micro segment.

The benefits of micro-segmentation include:

- Improves network access control to protect systems by limiting application access to a specific network segment or device.
- Happens at the application level (unlike firewalls) and can protect specific applications.
- Detects new or unsuspected network activity to and from a mainframe computer and blocks unauthorized users from connecting to an application. This approach ensures access only for authorized users and denies everyone else, a zero trust mandate.
- Reduces the potential risk should a network exposure occur.

Inherent mainframe characteristics make these goals difficult to achieve, however.

Traffic in and out of the z/OS mainframe uses Transmission Control Protocol/Internet Protocol (TCP/IP), which was designed to allow any-to-any connectivity with minimal configuration. This setup conflicts with security policies aimed at limiting connectivity to authorized users. The z/OS Communications Server includes controls in the System Authorization Facility (SAF), but the default for many sites is to allow all connections. TCP ports can be protected by SAF so that only permitted applications can open them, but furthermore complex controls are required to secure access to and from remote devices. Controlling tens of thousands of connection combinations can become an impossible task.

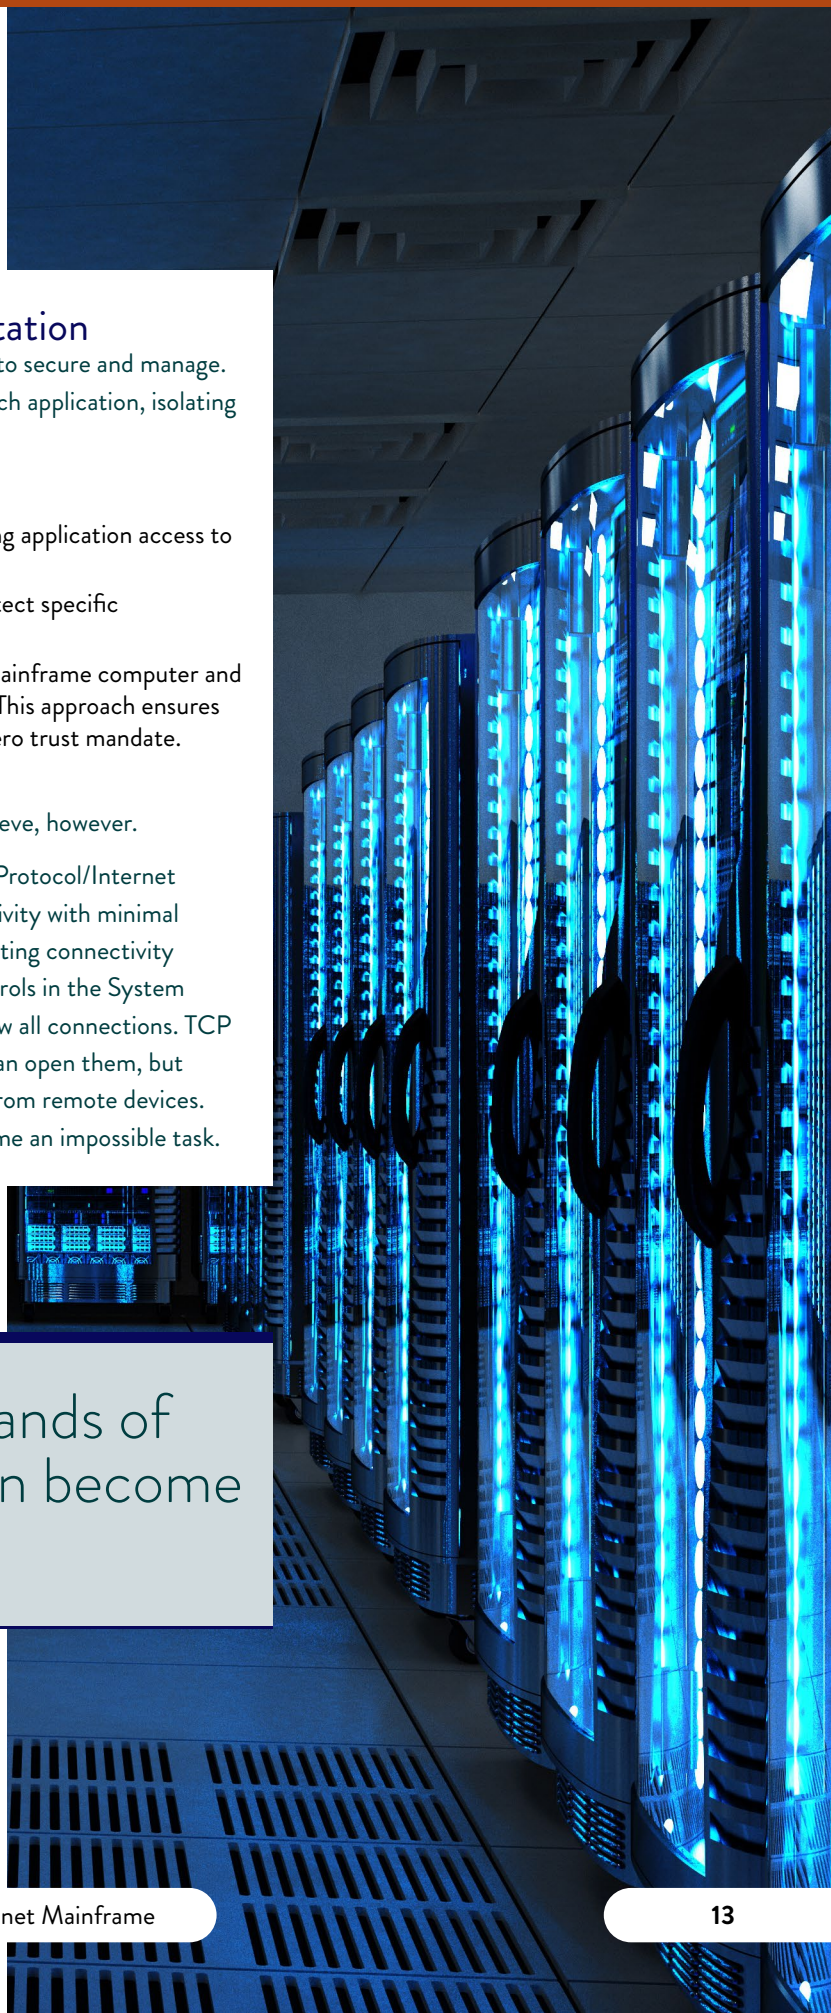

Controlling tens of thousands of connection combinations can become an impossible task.

Many mainframe sites lack an up-to-date and accurate picture of real-life network activity, such as which network devices are connected to specific applications and what is encrypted.

Most security mechanisms look at inbound TCP connections, but few look at controlling outbound connections. Any user can often initiate an outbound connection to a remote system, and hackers use outbound connections as a backdoor to mainframe services. User Datagram Protocol (UDP) activity is typically unsecured and unmanaged.

There are tools built into z/OS, but they can be difficult to configure and manage at large scale:

- IBM Policy Agent, part of Communications Server inside z/OS, can filter mainframe packets at the application level to provide segmentation but the process can be complicated, especially for organizations with thousands of connections.
- IBM z/OS Management Facility (z/OSMF) provides a graphical user interface (UI) that can be used to define policy agent filtering rules, but this requires time consuming manual data entry and knowledge of IP addresses and port numbers to add filters. As with Policy Agent, IBM z/OSMF does not easily scale for large organizations.

As a result, there is often a lack of understanding of what needs to be configured because application owners, network administrators, and security teams do not always have a complete picture. Among these groups, there can also be confusion over who is responsible for compliance such as PCI/DSS.

Dozens of applications running card data across hundreds of logical partitions (LPARs) could result in tens of thousands of network devices. All those devices become part of the PCI assessment scope unless card data applications can be segmented.

Solving these problems often requires a third-party tool that helps organizations understand what to configure, makes the configuration easy, and assigns configurations to the right group.

Vertali zTrust Manages Micro-Segmentation

Vertali zTrust for Networks manages micro-segmentation using IBM z/OS tools. Based on zTrust's network discovery capabilities, zTrust provides an understanding of network and traffic patterns, building a complete map of network connections to facilitate the micro-segmentation process. It works alongside controls managed by IBM z/OS such as user access, multifactor authentication (MFA) and encryption, providing a valuable additional layer of security.

zTrust gives security teams the ability to control access by permitting network segments to access applications through standard SAF controls and commands. It can detect new or unexpected network activity to and from the mainframe and confirm that the micro-segmentation settings are correct and working. zTrust automatically generates policy agent access control lists (ACLs) directly from SAF resources managed by standard External Security Manager (ESM) commands such as those provided with RACF, Access Control Facility 2 (ACF2) or Top Secret Security (TSS).

zTrust detects all traffic on an LPAR and builds a knowledge base of every mainframe connection. The first time zTrust detects an IP address connecting to an application, it records that in the knowledge base, together with the encryption status of that connection.

zTrust uses the knowledge base to build a complete set of External Security Manager (ESM) resources and access lists based on current network traffic. Security teams can review access lists to ensure only permitted network segments and devices are accessing key applications and access controls can limit access to encrypted network connections.

After analyzing the ESM profiles, zTrust builds IBM Policy Agent profiles that permit or block network traffic. zTrust makes segmentation simpler by managing ESM resources by name rather than IP addresses and port levels. It continuously monitors network activity to ensure the ESM policies defined are correctly implemented and to highlight any network changes that may require additional policies.

Over a short period of time, the knowledge base will provide a complete map of network activity by recording every unique connection. zTrust generates an alert when it detects an IP address connected to an application on the network for the first time. Filtering options are provided to whitelist resources to reduce alert volumes. zTrust alerts can be routed to offboard security information and event management (SIEM) solutions such as QRadar or Splunk via the Syslog Daemon.

zTrust documents all activity in audit logs and can generate periodic reports that confirm network micro-segmentation policies are implemented and a valuable resource to prove micro-segmentation is indeed in place and working.

zTrust also ensures connections are encrypted by differentiating between clear and encrypted network connections. It identifies applications that are permanently or temporarily accepting inbound non-encrypted or inbound encrypted connections and applications that are making outbound non-encrypted or outbound encrypted connections.

5 Stages of zTrust Software:

Stage 1

Network Discovery: A unique tool to build your network knowledge base and continuously monitor for new network activity.

Stage 2

ESM Resource Generation: Automatically generate ESM resource definitions and access lists for RACF, TSS or ACF2.

Stage 3

Build Security Profiles: Build policy agent profiles from ESM resources

Stage 4

Managed implementation of new policies with rollback option

Stage 5

Monitor and Manage Complexity: Monitor network activity and alert on policy violations

At any stage, reports can be produced to provide details on the SAF resources defined, permitted access lists for each application, the network connection maps and the live filters currently loaded into TCPIP.

Conclusion

Micro-segmentation makes it possible to logically divide networks into separate security segments at the level of specific workloads. By allowing organizations to define security controls and restrict access to each segment, micro-segmentation is an important step toward achieving Zero Trust. This security is crucial for financial institutions and others that hold sensitive customer information, often on mainframe computers.

Although micro-segmentation adds to the security of mainframe data, it is difficult to accomplish at scale. Large companies with thousands of network devices and applications might struggle to isolate all their resources without helpful third-party tools.

Vertali zTrust works by using standard IBM mainframe tools and interfaces. It adds management, implementation, and monitoring controls to isolate systems with different security needs. This approach reduces the number of systems in PCI DSS compliance scope and empowers the Cyber/Security teams to implement segmentation via their ESM. It also saves organizations time and money from performing these tasks manually.

zTrust blocks unwanted traffic and puts mainframe security where it belongs—in the hands of an organization's security team. It controls access by permitting network segments to access specific applications through standard SAF controls and commands. That provides micro-segmentation rather than blocking or enabling access to the entire mainframe.

Author Byline: This paper was written in partnership with The Futurum Group, an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day Futurum's analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

Click Here to Watch the Video