Technology Audit

Application Development

Azul Systems

Network Attached Processing / Azul Compute Appliance

Written by: Michael Azoff

Date: April 2006

Abstract

Azul Systems offers the first Network Attached Processing solutions for virtual machine applications. Named Compute Appliance, this device can currently hold a maximum 384 CPUs. Compute Appliances can be pooled to provide virtually unlimited processing capacity within a data centre, offering cost benefits in processing, power requirements, and capacity planning. Currently offering Java Virtual Machine (JVM) solutions, .NET solutions are in the pipeline. Azul proxy virtual machine software is installed on host servers that redirect traffic to the Compute Appliances. The cost benefit of Azul solutions, begin for data centres requiring 40 CPUs and upwards, but the company is first targeting the large, Global 2000 enterprises. Azul solutions allow server farms to scale with simple plug-and-play attachments of Compute Appliances, lower the cost of delivering Service Level Agreements, and make it easier to provide on-demand computing. Although latencies occur when applications need to interact with distributed systems, such as databases, these are at an acceptable level, and for computationally-intensive applications, particularly of a multi-threaded nature, there is a significant performance improvement to be gained.

KEY FINDINGS

Key: ✓ Product Strength ✗ Product Weakness *i* Point of Information

✓	Pioneering Network Attached Processing devices.	✓	Makes it easier to administer a server farm, and can use existing JVM monitoring tools.
✓	Ideal for heavy-compute applications and multi-threaded applications.	✓	Plug-and-play with no change necessary to the application.
✓	Pooled clusters of devices, each holding a current maximum of 384 CPUs, offering unbounded computing capabilities.	✓	As transaction loads increase, the Azul system shows superior throughput characteristics versus traditional servers.
✓	Simplifies capacity planning for individual applications.	✗	Currently available only for Java applications, although .NET is in the pipeline.

LOOK AHEAD

The second generation of the Vega chip will be released in 2007. There are also plans to release a Microsoft .NET solution. Butler Group expects Azul's novel Network Attached Processing solutions to create and grow a new market.

▶ FUNCTIONALITY

Product Analysis

Azul Systems has produced a series of Compute Appliances that represent the industry's first Network Attached Processing (NAP) solution. In what it expects will emulate the success of a similar concept, Network Attached Storage, NAP promises to transform the IT industry by providing plentiful processing capabilities at lower costs and lower power requirements. Compute Appliances deliver huge amounts of computational and memory resources as a shared network service, to applications built on the Java 2 platform Enterprise Edition (J2EE), with other Virtual Machine (VM) languages in the pipeline.

Compute Appliances are flat Symmetric Multiprocessing (SMP) systems with up to 384 processor cores and up to 256 GB of coherent memory in each appliance. However, the capacity of an Azul Compute Appliance goes well beyond traditional SMP systems, allowing many applications to share the same compute fabric with high levels of scalability, throughput, utilisation, and predictable response times under heavy load. This translates to a dramatic reduction in the number of servers required to build enterprise data centres, as well as the associated management costs.

The current Azul solution is 100% Java compatible and is designed to be operating system agnostic, meaning that any existing Java/J2EE application, regardless of operating environment, can access this shared physical resource, further reducing operating costs associated with traditional scale-out or scale-up deployment models. Adopting a NAP approach also dispenses with the need for individual application-level capacity planning. Furthermore Compute Appliances do not need to be configured and managed like traditional server architecture – expanding capacity is as simple as plugging additional appliances into the network.

Product Operation

At the core of the new technology is a segmented VM architecture. This is one of the critical pieces of intellectual property that Azul has pioneered – effectively separating 'compute' from the computer. This allows the redirection of VM-based applications, completely unmodified, to tap into the massively scalable compute capacity of Azul Compute Appliances.

Multiple appliances form compute pools, which are equipped with Compute Pool Manager, an easy-to-use Web-based monitoring tool that manages multiple compute appliances as a single-shared pool of processing and memory resources. Compute Pool Manager is able to reallocate CPU resources dynamically in real-time, on the order of 10 milliseconds, so quality of service can be guaranteed, helping enforce Service Level Agreements (SLAs).

Additionally Compute Pool Manager provides detailed data collection of resource usage for each application running in a compute pool – for each appliance in a pool and for each compute pool within a domain. The information gathered shows the identities of the monitored entities, along with current and cumulative measures of resource usage. This data can be used for a number of purposes, such as measurement and chargeback for resource usage, enhanced infrastructure cost analysis and budget planning, bill back models for usage, and tracking pool-level utilisation rates and trends around capacity planning. The gathering of metrics like these is a critical enabler to delivering NAP as a shared service, whilst a simplified billing process is a critical step on the path toward making utility computing a viable model for an enterprise IT infrastructure.

The underlying technology to NAP is, according to Azul, the industry's first and only multi-core chip designed for VM execution – the Azul Vega™ processor. The Vega chip has 24 cores, it does not expose its binaries to end-user applications, and is optimised specifically for VM workloads with features such as Pauseless Garbage Collection, Optimistic Thread Concurrency, and supporting up to 96GB heap sizes to dramatically improve the efficiency of running object-oriented programming languages. Azul can put up to 16 of these chips in its 384-way compute appliance, along with 256 GB of memory, which fits into an 11U rack, making it the world's largest SMP. This compute capacity has been built with high-levels of reliability, power, and floor space characteristics.

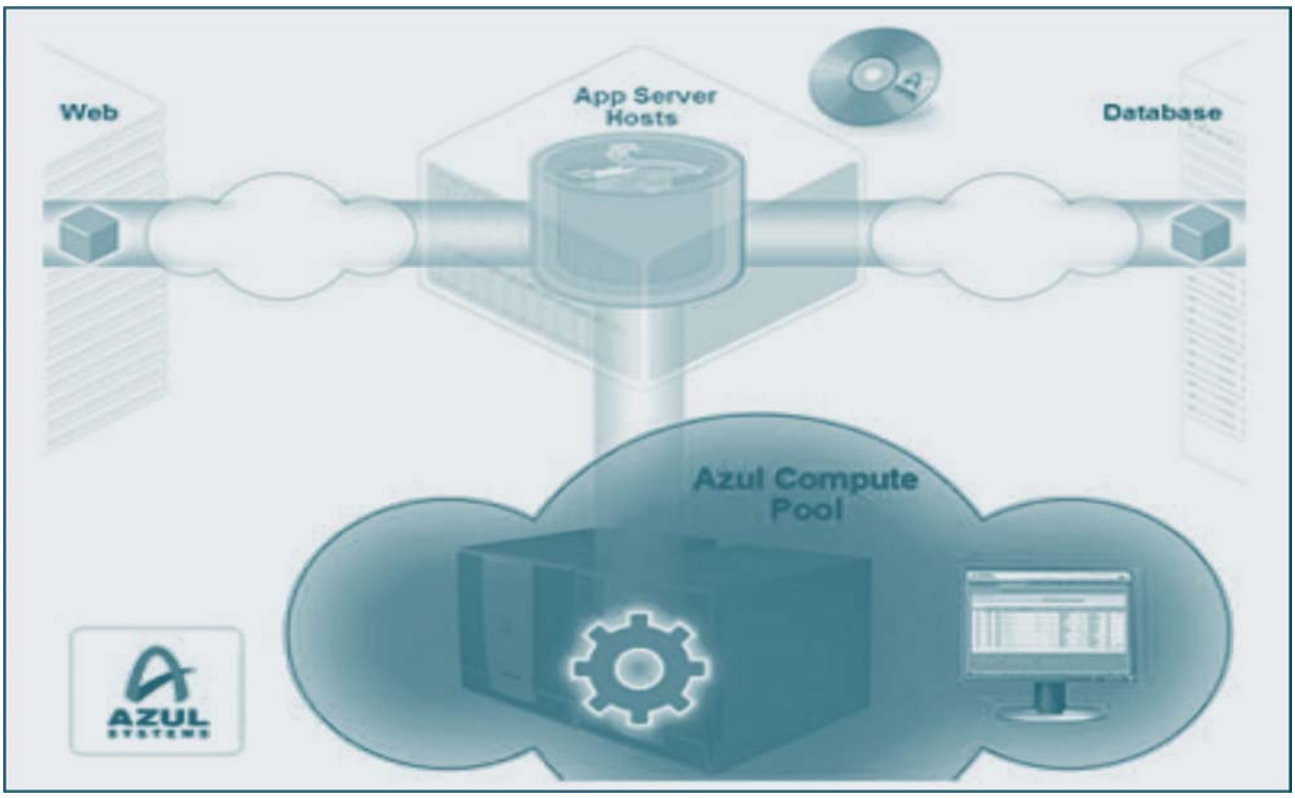

Figure 1: Network Overview of Compute Pool Installation

A Compute Appliance is 'plug and play': when fitted into an existing network infrastructure, the Compute Pool Manager discovers it and creates the pool. Azul VM proxy software is installed on the host system, establishing the redirection of the workload. From the user perspective the application appears to be running on a traditional server, but the proxy automatically transfers the application to the compute pool for execution. Figure 1 provides an overview of the Compute Appliance pool installation, and Figure 2 shows how the proxy software sits on the host server.

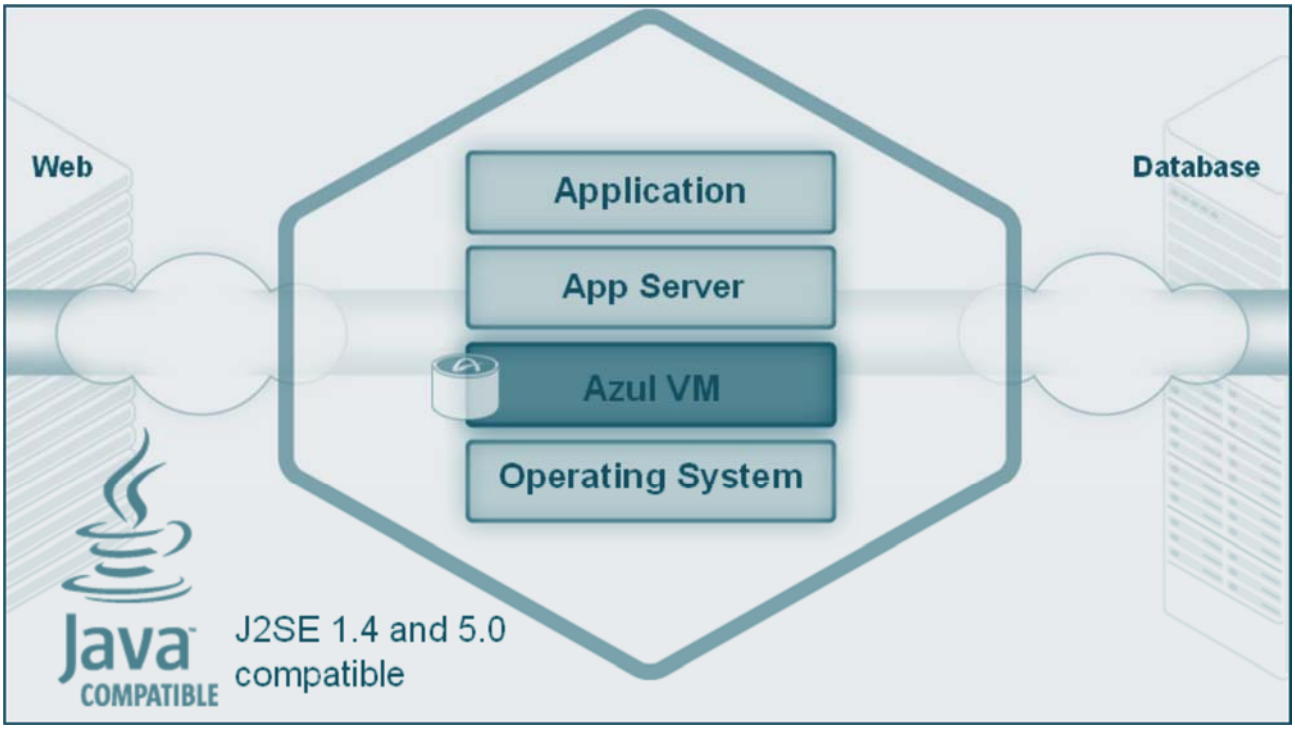

Figure 2: Position of Azul Virtual Machine

Each Azul Compute Appliance is designed to demonstrate extremely high levels of fault tolerance. The appliances are designed to have no single point-of-failure, and are fitted with dual-redundant network processor cards, each with dual gigabit Ethernet connectivity, N+1 power supplies, and N+1 cooling. Both fans and power supplies are field replaceable units. The rest of the appliance is entirely solid state, and so it is highly reliable. There is parity checking on all processor caches and registers, with ECC (chipkill) on all the main memory. Should components fail, appliances can reboot in a partially de-configured state, which maintains provision of some level of capacity prior to service personnel coming on-site to replace the failed unit.

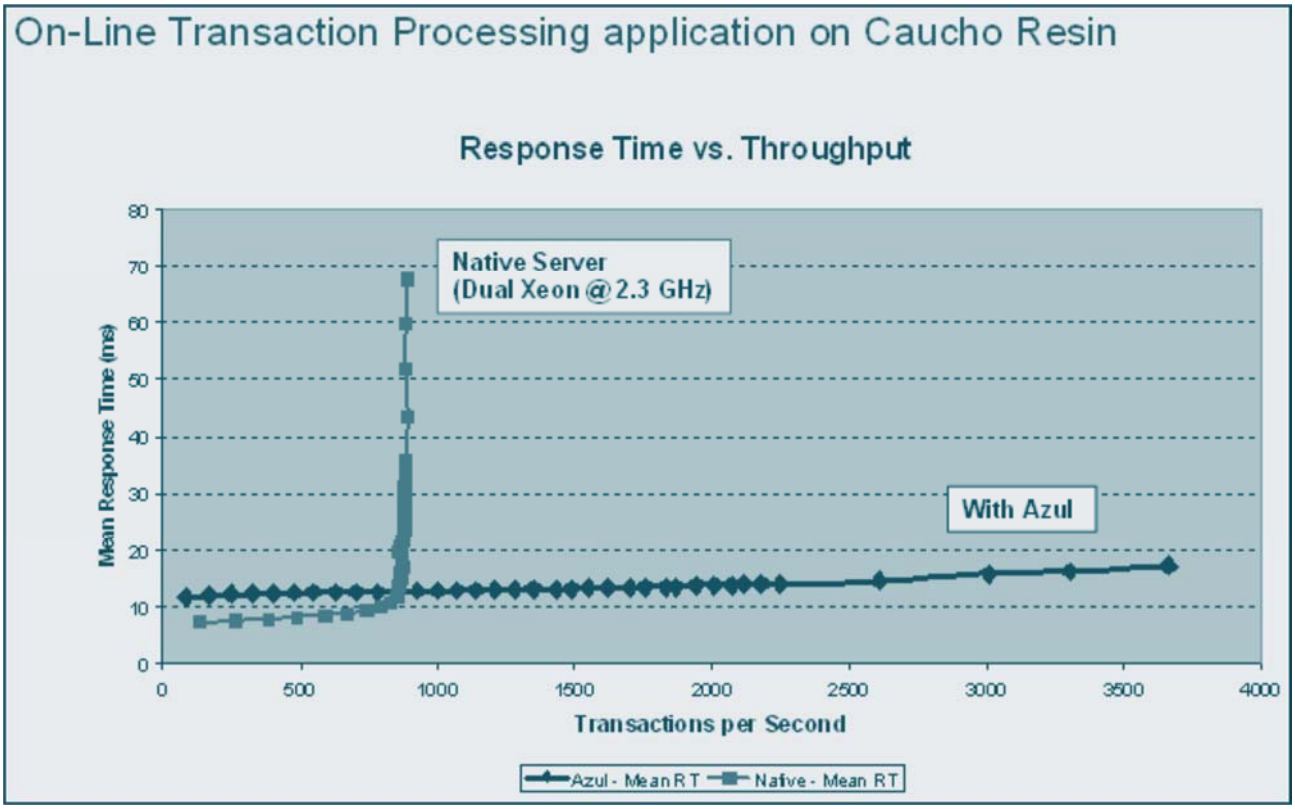

Figure 3: Response Time Versus Throughput Comparison for a Traditional and Azul System

The design Mean Time Before Failure (MTBF) of an individual appliance is 40,000 hours and Azul has empirically demonstrated an MTBF of at least 11,000 hours with over 90% confidence. The MTBF of a whole pool is much greater than that of a single Compute Appliance, and the Mean Time to Recovery for an individual VM instance is very low.

Java application transactions streaming through the proxy server incur an additional network hop in the journey to and from the Compute Appliance. This latency, of the order of milliseconds, may be visible under low transaction volumes, compared with traditional systems, however, under typical server farm duty loads, the removal of processing bottlenecks shows NAP to be superior in transaction response times. Figure 3 shows the characteristics of a traditional Application Server, the Caucho Resin J2EE Application Server, compared with a NAP system: under a heavy workload there is a sharp degradation in the former, while the Azul system shows a steady performance under increasing load.

In practice, multi-tier applications involve multiple network hops, diminishing any sensitivity to the hop to NAP. Communication between the Proxy and the Compute Appliances also employs caching to minimise the impact of network latencies. Advanced network caching technologies from companies such as Gigaspaces, Tangosol, and Terracotta also complement NAP. In the future, the deployment of layer-4 and layer-7 switches, promise to further optimise network traffic across a distributed application. These emerging technologies will further highlight the value of NAP.

Product Emphasis

Azul stresses it is not a chip company contrary to first impressions. Its intellectual property covers managing multiple JVMs running multi-threaded applications, and offers Pauseless Garbage Collection. Garbage Collection is the automated management of objects in the JVM. In traditional large-scale Java deployments, applications need to be paused for Garbage Collection to catch up. Azul's pauseless advance has Garbage Collection running concurrently with applications, and with multiple collectors running in parallel. Heap de-fragmentation is also performed concurrently. If a better multi-core chip appeared on the market, Azul says that it can simply swap its software products to the new device.

▶ DEPLOYMENT

In order to move applications over to an Azul compute pool, it is simply a matter of replacing whatever conventional JVM it was running on with the one supplied by Azul – depending on a customer's operating procedures, this typically involves altering application server start scripts (so that Java applications are pointed to the Azul Java Developer Kit) and regression testing the application in the new environment. This work would be carried out by application engineers or infrastructure operators. A full production rollout of Azul would typically take as long as it would take an organisation to move a given Java application to a new version of the JVM. This will range from 15 to 60 elapsed working days.

In a production setting, it is also likely that the compute pool will be integrated into an existing Simple Network Management Protocol (SNMP) management infrastructure. Azul appliances issue SNMP traps to alert administrators to significant events, and also offers browser accessed Management Information Base, to enable monitoring from SNMP-enabled enterprise consoles.

The product is designed to be deployed into a compute pool in an incremental manner. One of the purposes of pooling capacity in this way is to allow capacity to be built up on-line on an as-needed basis – utilisation can be tracked at the pool level, and when some customer defined comfort threshold is passed, further appliances can be purchased and added to the compute pool with no need for change at the application tier. This is very much analogous to the simplified, transparent way that storage is added to a Network Attached Storage deployment.

The compute pool administrator sets policies for all applications tapping into the compute pool. Policies can be very granular, but in general, they are set minimum and maximum limits on resource utilisation and availability configurations. One of the key strengths of NAP is to significantly reduce the management overhead for Java application deployments. This is achieved by reducing the overall number of servers required to run Java applications and by providing a centrally-managed large pool of capacity. Thus, the need for personnel to administrate the compute pool is more than offset by the reduced headcount brought about by reducing the number of UNIX boxes needed to run the production workloads. This fact is core to the Azul Total Cost of Ownership (TCO) model, which indicates potential savings of 50% over a three-year period of operation, with the payback typically happening well within the first year of operation.

Training is supplied on-site as part of the initial installation package. There is also a Quickstart package: a Web-based introduction on how to set up and manage Azul Compute Appliances, application management, and using the Compute Pool Manager tool. Instructor-led classroom and lab training is provided through Azul's Compute Pool Administration Training, which teaches the essential fundamentals of administrating and operating Azul NAP solutions.

Further educational training can be provided to Java architects and developers as to how their programming style can now change since NAP removes the need to impact applications that had originally been designed for traditional VM environments – for instance, applications no longer need be split into segments to avoid large Java heaps which would have caused unacceptable pauses for Garbage Collection; also, less complex locking strategies can be adopted due to the hardware-supported Optimistic Thread Concurrency.

Ongoing support includes software and hardware support, round-the-clock call-centre availability and a support portal in the event of any problems. For production customers, support includes a four-hour on-site parts replacement, and "follow the Sun" telephone and Web support, which are handled out of Azul's global support centre in California. Azul also has regional centres in the UK and Bangalore. A license subscription provides access to software updates and new platform ports as they become available.

Azul's remote support monitoring can automatically detect installed Compute Appliances and transfer issues as they arise to the Azul Customer Advocacy Centre as required. If a failure does occur, an Azul-certified technician is dispatched along with the replacement parts to the customer site to ensure that any downtime is minimised.

Azul supports a range of platforms including BEA WebLogic, IBM WebSphere, Caucho Resin, JBoss, and Apache Tomcat. It has proxy solutions for Solaris, Linux, HP-UX, and AIX operating systems. Compute Appliances have been designed to support the Microsoft .NET Framework, but an announcement on when proxy support will be available, has not yet been made.

Any Java 2 platform Standard Edition 1.4 and 5.0 application is eligible to run on Azul Compute Appliances. If these applications contain elements of legacy connectivity and integration logic, then this will continue to function unchanged following the introduction of Azul.

Model Number	Model 960	Model 1920	Model 3840
No. of Vega chips per appliance	4 chips	8 chips	16 chips
No. of processor cores per appliance	96 processor cores	192 processor cores	384 processor cores
Memory	16 or 32 GB memory	32 or 64 GB memory	128 or 256 GB memory
Power consumption	950w	1600w	2700w

Azul products are available in multiple configurations:

Software licenses are included with the purchase of the appliances, and each appliance is sold in conjunction with Azul "Pool Power" – a service product that includes hardware maintenance, software maintenance, and support, and also includes 24x7 telephone support. Pricing starts at UK£70,000 for a 960 appliance (with 96 CPU cores and 32GB of memory, primarily used for development environments) and ranges up to UK£420,000 for a 3840 appliance (with 384 CPU cores and 256GB of memory) being deployed in high availability production environments.

▶ PRODUCT STRATEGY

Azul's target audience includes both vertically and horizontally focused enterprises. It began shipping Compute Appliances in late summer 2005 and is working with enterprise data centres across many sectors, including financial services, telecommunications, e-commerce, hospitality, healthcare, and central government, to name a few.

To help customers understand the benefits of NAP under different scenarios, Azul has developed a TCO calculator that compares the TCO and Return On Investment (ROI) of different application deployment options. This calculator was developed in collaboration with IT organisations, to analyse the costs associated with application deployment. While the Azul solution typically complements existing scale-out or scale-up deployments, it is often compared to these models. As such, the TCO calculator, models the costs of deploying systems in the application tier via scale-out, scale-up, or NAP, allowing cost comparisons and "what if" analyses of all three.

The TCO model takes into account such factors as acquisition costs, including the number of servers deployed, price per server (warranties, pre-paid maintenance, support, and operating system and right-to-use licenses), annual administration costs (staff costs divided into two main subcategories: server administration and application administration), and annual facility costs (based on power, cooling, and data centre facilities). These elements of the overall "lights on" costs, provide a relatively accurate TCO, with little room for subjective measures to skew the figures and impact the credibility of the analysis.

The initial target companies for Azul Systems are the Global 2000 where larger enterprises will see the best returns on data centre consolidation, although the starting point can be as little as 40 CPUs of Java (in development, production, or both), which is the minimum CPU count that would be required to make the economic case for NAP. Azul's go-to-market strategy is a direct sales model.

Azul has gained significant traction with many world-class alliances, partners, and ISVs who share in its vision, and will continue to drive innovation and maturation around the model of NAP. Azul has signed partnerships with BEA and JBoss, which include agreed upon licensing and certification agreements. Additionally, Azul is actively working with IBM and Oracle to further formal certification and support relationships. Azul has now verified more than 120 commercial ISV applications that are suited to run on the company's new compute platform. Also, relationships are building with IT outsourcing players, such as EDS, who views Azul as an opportunity to increase the service levels, whilst also lowering the operational cost of its outsourced data centre operations, and significantly increasing its operating margins.

Azul Systems also offers premium customer support through IBM Global Services, around the world. The company is also building up its relationships with systems integrators in Japan and in the UK to extend its reach into these markets.

Azul has a regular release cycle for minor upgrades (two to three per year) and a major release is planned for 2007 (the Vega 2 chip). A key element of Azul's release strategy lies in the fact that its systems do not expose the Azul instruction set and rely on the Java layer for compatibility. This ensures the release strategy is as non-disruptive as possible.

▶ COMPANY PROFILE

Azul Systems is a privately-held company, headquartered in Mountain View, CA, with regional offices in Japan and the UK, and has a development centre in India. It was founded in 2001 by Gil Tene, Shyam Pillalamarri, and Scott Sellers. It is funded by well-respected venture capitalists firms, including Accel Partners, Austin Ventures, comventures, redpoint Ventures, and Worldview Technology Partners. Credit Suisse, initially a customer, has also taken an investment position in the company.

The company was formed with the specific aim of creating a new IT architecture exploiting current approaches in application development, such as VMs, and has named this architecture Network Attached Processing (NAP). The founders of the company wanted to match the hardware infrastructure to the applications that are being developed within VM-based languages, and in doing so shift this to a hardware platform that was more suitable to supporting these applications. This approach mirrors the way that other capabilities that previously existed on the server, such as routing and storage, have been offloaded to specialised devices. With NAP, computing capability has become a resource to be managed and delivered to the organisation in the same way as storage.

The company currently has 200 employees in engineering, sales, marketing, customer support, and administration.

Azul has opened the Centre for Unbound Compute in Mountain View, where registered users can evaluate the advantages that NAP offers to Java applications. The company also has an Azul performance centre in Tokyo through its partnership with NISSHO, one of the company's three Japan-based resellers.

Azul has evaluations at close to 50 enterprise data-centre sites.

▶ SUMMARY

As enterprises roll-out new applications and services, grow their user base, and tax their data centre infrastructure with voluminous workloads, the benefits of existing scaling models (scale-up/out) are no longer delivering the value that they once did. Traditional systems architectures have not been designed for application VMs, or the highly multi-threaded and "Garbage Collected" nature of the applications that utilise them. As a result, a widening gap has emerged between the compute infrastructures that today's applications can exploit and the reality of what today's data centre can deliver. It is this gap that provides the opportunity for NAP.

NAP offers: improved efficiency and utilisation through shared computing resources; continued isolation of applications from each other; a server consolidation solution, leading to reduction in power and footprint; reduced capacity planning effort, since compute power is centralised; reduced TCO and simplified administration management; the ability to respond quickly to peak demands, and to move towards a "capacity on demand" model, with corresponding business benefits; and consistent service levels to end-users, irrespective of the loading. NAP can help realise the delivery of IT as a service, and is in line with current trends towards Service Oriented Architecture. Butler Group believes that Azul Systems offers innovative solutions for data centres that deliver substantial benefits.

Contact Details

Azul Systems, Inc.
1600 Plymouth Street
Mountain View
CA 94043
USA

Tel: +1 650 230 6500
Fax: +1 650 230 6600

E-mail: info@azulsystems.com

www.azulsystems.com

Azul Systems UK Limited
Regus House, 268 Bath Road
Slough
Berkshire, SL1 4DX
UK

Tel: +44 (0)1753 708570

E-mail: azuleurope@azulsystems.com

Butler Group
a Datamonitor Company

Headquarters:

Europa House,
184 Ferensway,
Hull, East Yorkshire,
HU1 3UT, UK

Tel: +44 (0)1482 586149
Fax: +44 (0)1482 323577

Australian Sales Office:

Butler Direct Pty Ltd.,
Level 46, Citigroup Building,
2 Park Street, Sydney,
NSW, 2000, Australia

Tel: + 61 (02) 8705 6960
Fax: + 61 (02) 8705 6961

End-user Sales Office (USA):

Butler Group,
245 Fifth Avenue, 4th Floor,
New York, NY 10016,
USA

Tel: +1 212 652 5302
Fax: +1 212 202 4684

For more information on Butler Group's Subscription Services please contact one of the local offices above.