

July 2010

VALUE PROPOSITION FOR IBM zENTERPRISE SYSTEM: NEW DIMENSIONS IN WORKLOAD MANAGEMENT

Platform

New Opportunities

IBM's new zEnterprise System represents a major transition in mainframe architecture. Established System z strengths in transaction processing, data management and movement, availability, security and manageability are reinforced. At the same time, changes in processor technology, memory structure and overall system architecture break new ground in a number of areas.

This report focuses on one of these areas – workload management. The new zEnterprise Unified Resource Manager extends System z management strengths to selected IBM blade servers. The result is to enable extremely granular, real-time management of “end-to-end” workloads across multiple server platforms and architectural tiers.

There are, in many large organizations, applications for which this approach makes sense. Complex architectures typically include at least three server tiers – back-end systems and data, application servers and Web servers. To these must be added servers supporting firewalls, directory services, encryption, caching, load balancing, management and a variety of other functions.

With such architectures, performance and integrity may become increasingly problematic. Challenges are compounded when extremely large workloads and high throughput volumes must be supported. Even if service-level targets can be adequately met, costs of operation tend to escalate.

zEnterprise systems offer a new way of addressing these challenges. By combining physical and logical server consolidation with tightly integrated, highly granular, real-time workload management, it becomes possible to achieve new levels of end-to-end performance and quality of service.

Where mainframes constitute the back-end data tier of complex business-critical server architectures, zEnterprise systems may offer new, cost-effective opportunities to extend “mainframe-class” capabilities across all platforms and tiers. Three examples of such opportunities – complex Internet architectures, large ERP systems and organizational data warehouses – are outlined in this report.

Hybrid Architecture

Workload management forms part of a larger set of zEnterprise capabilities announced by IBM. In the new-generation System z196, for example, changes in processor and memory structure, in compiler technologies and in other components offer the potential for significant improvements in performance and scalability within the z/OS environment.

zEnterprise introductions continue the evolution of the System z toward a hybrid architecture incorporating diverse specialized engines. This began with earlier Application Assist Processors (zAAPs), Integrated Information Processors (zIIPs) and other specialty processors. zEnterprise represents a major new phase in this evolution.

zEnterprise systems combine z196 frames with IBM BladeCenter frames equipped with POWER7 and/or System x (x86-based) blade servers running AIX or Linux operating systems respectively. Support for System x blades, according to IBM, will be in the first half of 2011.

The IBM Smart Analytics Optimizer (SAO), a blade-based appliance that accelerates certain types of high-volume analytical query, is also supported.

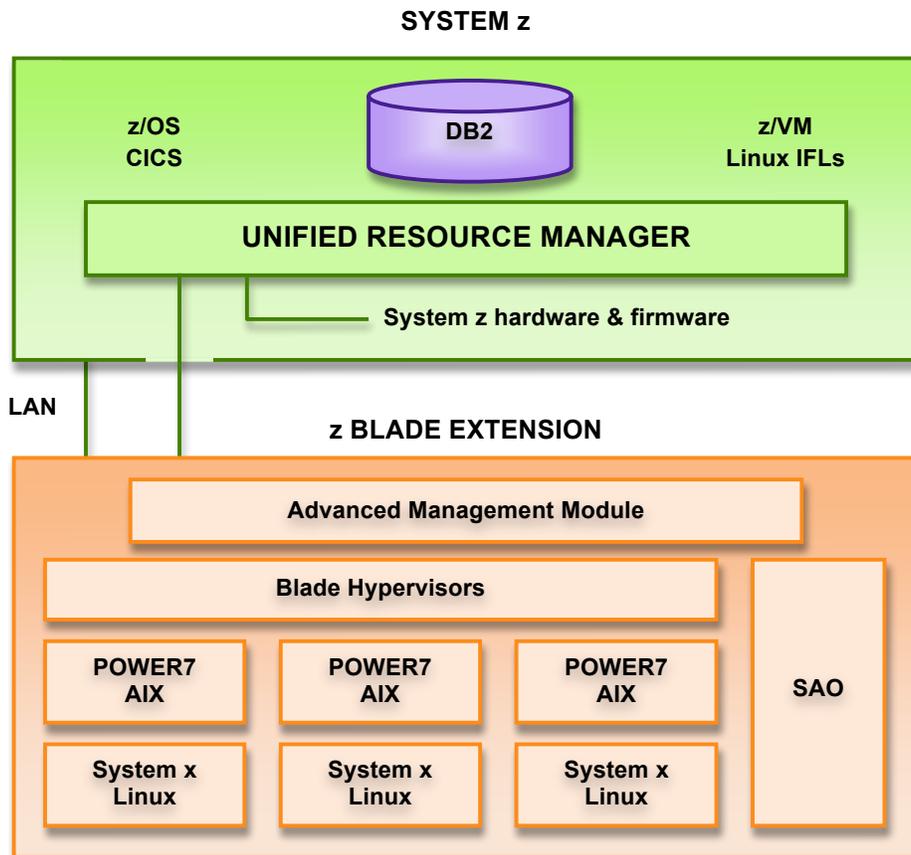
The move to hybrid architecture reflects industry trends. In most organizations, workloads have grown progressively more complex as distinctions between transactional, analytical, collaborative and other systems have eroded. Consolidation of servers, as well as closer integration of diverse software components and pressures for “real-time” interaction between these have contributed to this process.

In such environments, it may no longer be appropriate to define performance in application- or workload-specific terms. The key variable becomes the efficiency with which systems execute heterogeneous workloads with high levels of interaction and concurrency. Undifferentiated single-processor architectures may no longer represent the most efficient approach.

System efficiency, in turn, is determined by workload management mechanisms. The effectiveness with which these allocate and re-allocate physical as well as virtual resources materially affects system-wide performance, capacity utilization and the extent to which outages and bottlenecks may occur.

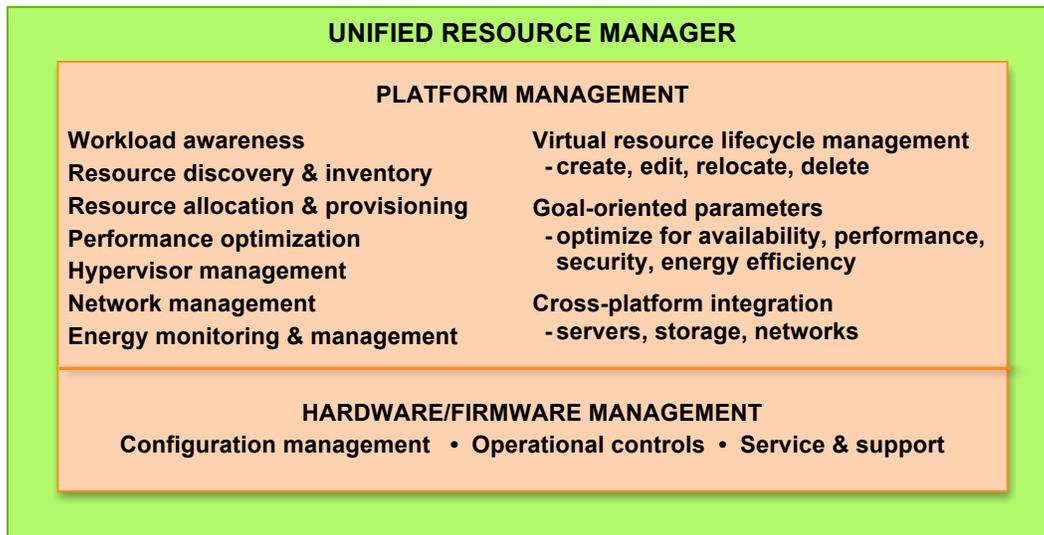
The zEnterprise system thus represents more than a “blade chassis bolted to a mainframe.” A sophisticated workload management structure, illustrated in figure 1, links the two platforms.

Figure 1
zEnterprise Workload Management Structure



The Unified Resource Manager delivers the services shown in figure 2.

Figure 2
Unified Resource Manager Services



The Unified Resource Manager integrates management of virtual as well as physical resources. It links directly to z/OS and z/VM hypervisors, as well as to PowerVM hypervisors on POWER7 blades. Support for the kvm Linux hypervisor, according to IBM, will be in the first half of 2011.

The Unified Resource Manager provides virtual lifecycle management for all supported hypervisors, allowing virtual resources be created, modified, moved and deleted as required. Comparable capabilities are provided by IBM Systems Director solutions for Power and System x servers which are not integrated in zBladeCenter Extension (zBX) frames, and for IBM distributed storage systems.

zEnterprise workload management capabilities are derived primarily from the z/OS environment. The Unified Resource Manager, in particular, is largely based on the z/OS Workload Manager (WLM). This is a central component of the zEnterprise value proposition.

For decades, mainframes have been industry leaders in mixed workload management. Because they have typically run dozens to hundreds of diverse applications, mechanisms for allocation and re-allocation of system resources have become highly developed. Prioritization, scheduling, management of concurrent execution and other processes operate in an extremely granular manner, in real time.

Workload management functions, moreover, are closely integrated and optimized across the major z/OS software subsystems, including DB2, CICS and WebSphere, and extend to z/OS and z/VM virtualization technologies. They also dovetail with functions that provide availability, security and disaster recovery, in which mainframe systems are, again, recognized industry leaders.

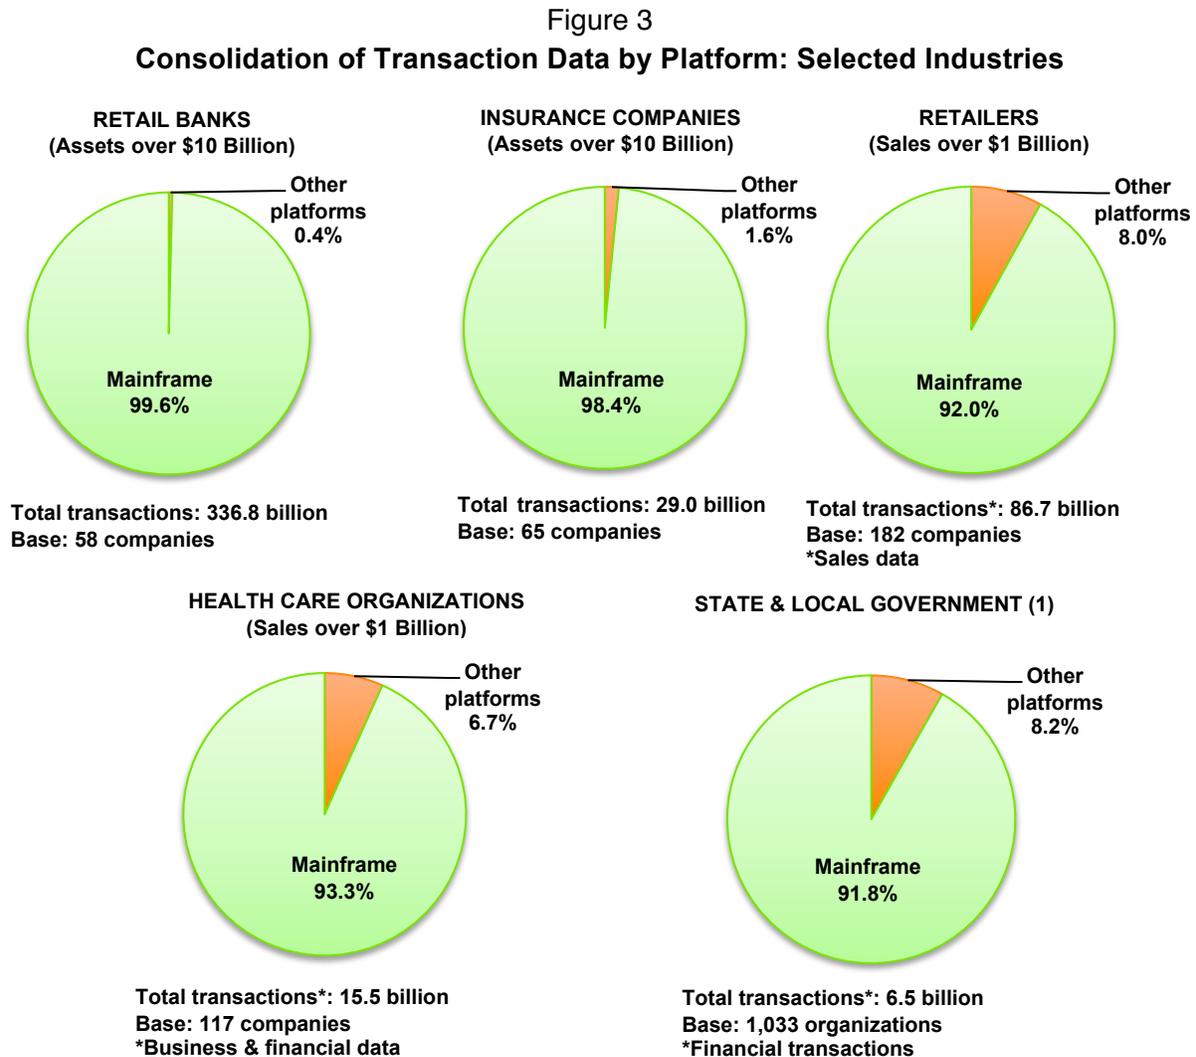
The goal of IBM zEnterprise workload management strategy appears to be clear: to extend these capabilities to all AIX and Linux system resources within zEnterprise systems.

Applications

Complex Internet Architectures

In many large organizations, mainframes continue to host core business-critical systems. In industries such as banking and insurance, mainframes are de facto global standards for core transaction processing. In others, they continue to handle very large transactional systems and support legacy applications that it would not be feasible or cost-effective to move to other platforms.

Where mainframes play such roles, they have typically emerged as the main point of consolidation for transactional data. Figure 3 shows examples.



(1) States, counties & municipalities with 100,000+ population; school districts with 100,000+ students

Source: International Technology Group

Although large segments of application and Web processing have moved to distributed platforms, transactions have continued to flow to and from mainframe systems. Transaction records are maintained and updated in mainframe databases. This has continued to be the case where organizations have built out multi-tier architectures to interact with customers, suppliers and business partners through the Internet.

In principle, most such architectures are three-tier. In practice, however, they have become structurally more complex. Layers of firewalls protect each tier, and specialized servers for directory services, load balancing, encryption and other functions have been incorporated. Key servers in each tier are clustered, and switches and routers are duplexed for availability purposes.

These structures can pose major challenges, particularly if very large transaction volumes must be supported. Maintaining response time may be difficult – a transaction initiated via the Internet routinely passes through 20 to 50 devices before it reaches a back-end database, and the number may be significantly larger – and availability, security and recoverability may become problematic.

In such structures, effective workload management can materially improve the efficiency with which server resources are used. This has two effects: throughput is accelerated, and risks that workload surges will overload systems are reduced. Outages and performance bottlenecks become less likely.

The zEnterprise approach has drawn the attention of a number of large players in financial services and other industries where customer bases have expanded rapidly after mergers and acquisitions. These and other organizations that have experienced sustained growth in Internet transaction volumes have often found that their architectures must deal with workloads a great deal larger than was originally expected.

There has also been a growing recognition that systems that interface to customers via the Internet have become as “business-critical” as conventional forms of transaction processing. For example, the bottom-line impact of Internet banking disruptions can be as great as outages affecting ATM or branch networks.

There is also a great deal of evidence that Internet customers who experience outages, slow response times or impaired applications are more likely to defect. Even if defections cannot be tied to specific experiences, overall satisfaction levels may be reduced – and decades of experience have shown that these materially affect attrition rates.

Clearly, the requirement is for “mainframe-class” quality of service for critical Internet systems. Few organizations would go as far as to implement entire multi-tier architectures on mainframes. But zEnterprise systems offer an alternative.

ERP Systems

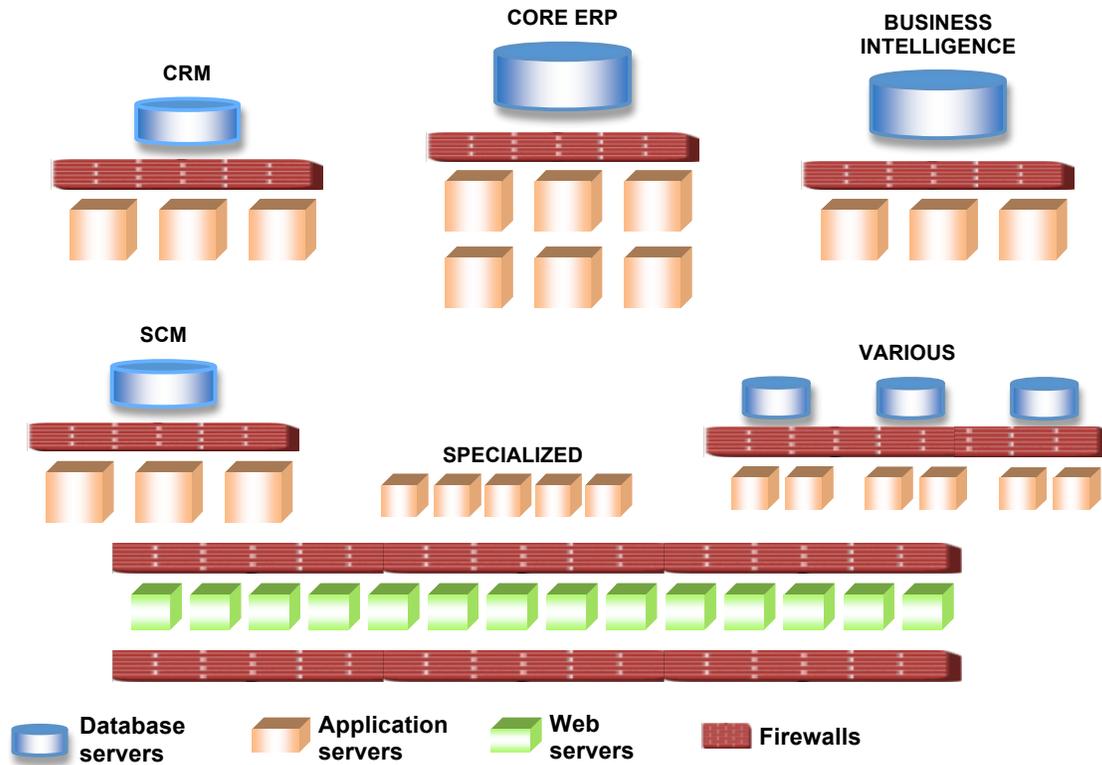
Most enterprise resource planning (ERP) systems are also built around multi-tier architectures. Where the database component of these is mainframe-based, UNIX and/or x86 servers are normally employed as application, Web, firewall, and specialized servers. In large installations, structures may become as complex as those for Internet architectures.

Several trends, however, have caused users to reevaluate such approaches. One is that workloads have not only expanded, but also become increasingly diverse. While early deployments focused primarily on transactional processes, organizations have progressively added customer relationship (CRM), business intelligence (BI), supply chain management (SCM), e-commerce and other types of system.

Latest-generation SAP industry solutions, for example, now routinely address between 300 and 500 separate business processes and functions. The workloads generated by these may include heterogeneous transactional and query processes; analytical computation; interaction with employees, customers and partners through Internet and intranets; data and document management; and others.

Where the different components of the SAP environment are deployed on physically separate servers, data paths may become highly complex. Figure 4 shows an example.

Figure 4
Complex SAP Installation: Example



A second trend has been acceleration of business operations. Cycle times have been cut, inventories reduced and processes streamlined. In manufacturing, retail, distribution and other supply chain-intensive industries, adoption of “lean” strategies employing just-in-time techniques become has become the norm. Other industries are moving in the same direction.

Acceleration magnifies stresses on underlying system infrastructures. Interactions between different modules and systems occur more frequently, increasingly heterogeneous workloads must be executed concurrently rather than sequentially, and replication processes become continuous rather than periodic. Highly effective workload management is again mandated.

The combination of workload diversity, growth and acceleration has caused the same performance and availability challenges to emerge as in complex Internet architectures. The potential impact of outages and bottlenecks may be equally severe.

Experience has shown that, in tightly integrated supply chains, the effects of disruption may “cascade” rapidly across organizations. A wide range of internal processes may be impaired, the effects may spread to customers, suppliers and partners, and dislocation may continue long after service has been restored. Where businesses operate in “real time,” with few or no inventory buffers, impacts are greater.

In organizations where mainframes act as SAP database servers, zEnterprise systems may be an obvious choice to improve system-wide performance and availability.

Data Warehousing

The term “data warehouse” is typically associated with standalone systems built around UNIX or x86 servers, or specialized platforms from vendors such as Teradata and Netezza. Oracle has also targeted this market with its Database Machine.

The System z, however, has been widely employed in some industries as a host for what can be characterized as “general purpose” data warehouses. Specialized systems were originally designed to serve small communities of users initiating high-volume analytical queries.

Increasingly, however, organizational data warehouses have come to support larger populations of managers, professionals and front-line employees whose informational requirements vary widely. Applications range from simple ad hoc query and reporting, to large-scale data scans and online analytical processing (OLAP) jobs.

There has also been a widespread trend toward use of “operational BI,” in which data warehouse access is provided to call center operators, customer service personnel and others handling large numbers of time-sensitive queries.

In such environments, systems must handle highly diverse workloads in an efficient manner. Concurrent processing becomes the norm, and granular scheduling is required to ensure prioritization. A large resource-intensive job initiated by one individual, for example, cannot be allowed to delay response to more important time-sensitive queries.

Challenges are compounded by acceleration of refresh cycles. Across a wide range of industries, monthly update cycles have moved to weekly, weekly cycles have moved to daily, and refreshes of data warehouse contents now commonly occur hourly, or on continuous basis. Data movement and query processes must, in a growing number of organizations, be conducted concurrently.

A familiar picture emerges: highly effective workload management is mandated.

The attractions of employing mainframes to host general-purpose data warehouses have been reinforced by other factors. Where core transactional databases also run on mainframes, synergies in data structures may be exploited, and extract, transformation and load (ETL) cycles may be shortened. The high-bandwidth, high-volume I/O strengths of mainframe architecture also favor this type of system.

With zEnterprise, the scope of mainframe-based data warehouses may be extended to handle a broader range of applications, while resource-intensive queries may be offloaded to the Smart Analytics Optimizer, or to POWER7 or System x blades. The capabilities of the Unified Resource Manager will enable the overall portfolio of query workloads to be executed with a high degree of efficiency.

Conclusions

Two main conclusions may be drawn.

First, where organizations already have significant investments in mainframes, and run core business-critical transaction processing systems on these, there are a number of distinct application areas where zEnterprise may provide value. A new fusion of mainframe, UNIX and x86 capabilities is enabled.

Mainframe strengths in workload management and quality of service may be extended to POWER7 and System x blade servers running AIX and Linux operating systems, and employing PowerVM and kvm hypervisors. IBM Smart Analytics Optimizers may also be exploited. Second, the full extent of this value may not yet be apparent. The IBM July 2010 announcements clearly represent early and partial implementations of zEnterprise architecture.

Further advances are likely in application and workload performance acceleration, using specialized processors and software optimization.

POWER7 processors also incorporate hybrid features enabling acceleration of query- and transaction-intensive workloads (e.g., TurboCore and MaxCore modes), and the overall POWER7 architecture appears to be moving in this direction. Such capabilities may be further leveraged in zEnterprise systems.

The role of firmware (i.e., chip-embedded code, also referred to as microcode) in zEnterprise systems is notable. The Unified Resource Manager is implemented in firmware, and this is also the case for the BladeCenter Advanced Management Module, with which it interfaces, and for PowerVM hypervisors. According to IBM, kvm hypervisors will also be implemented in this manner.

As a general principle, firmware generates lower system overhead than software, and is appropriate for real-time system processes requiring extremely low levels of latency. Use of firmware in zEnterprise systems, however, may evolve further, and may have significant architectural implications.

The capabilities announced by IBM in July 2010 may thus turn out to be the “tip of the iceberg” in terms both of zEnterprise architecture, and of the role of hybrid computing models in IBM product strategy.

Stay tuned.



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

Copyright © 2010 by the International Technology Group. All rights reserved. Material, in whole or part, contained in this document may not be reproduced or distributed by any means or in any form, including original, without the prior written permission of the International Technology Group (ITG). Information has been obtained from sources assumed to be reliable and reflects conclusions at the time. This document was developed with International Business Machines Corporation (IBM) funding. Although the document may utilize publicly available material from various sources, including IBM, it does not necessarily reflect the positions of such sources on the issues addressed in this document. Material contained and conclusions presented in this document are subject to change without notice. All warranties as to the accuracy, completeness or adequacy of such material are disclaimed. There shall be no liability for errors, omissions or inadequacies in the material contained in this document or for interpretations thereof. Trademarks included in this document are the property of their respective owners.