

Enterprise IT Trends and Implications for Architecture Research

Parthasarathy Ranganathan and Norman Jouppi

ISSL/MMSL

Hewlett Packard Labs

Palo Alto, USA

Email: Partha.ranganathan@hp.com, norm.jouppi@hp.com

ABSTRACT

The last decade has seen several changes in the structure and emphasis of enterprise IT systems. Specific infrastructure trends have included the emergence of large consolidated data centers, the adoption of virtualization and modularization, and an increased commoditization of hardware. At the application level, both the workload mix and usage patterns have evolved to an increased emphasis on service-centric computing and SLA-driven performance tuning. These, often dramatic, changes in the enterprise IT landscape motivate equivalent changes in the emphasis of architecture research. In this paper, we summarize some recent trends in enterprise IT systems and discuss the implications for architecture research, suggesting some high-level challenges and open questions for the community to address.

INTRODUCTION

The last decade has seen several changes in the structure and emphasis of enterprise IT systems. These have been driven by the growing adoption of the internet, the increasingly closer tie-ins between the IT infrastructure of the enterprise and its overall business processes, and the greater emphasis on the costs associated with managing IT. For example, most companies, even traditionally non-IT companies, now have vibrant e-commerce sites and well developed online business-to-business channels or customer-relationship portals that are closely tied to the success of their products. The capital and recurring costs needed to build and manage the infrastructure to power these services can often be a significant fraction of the IT operational costs an enterprise is trying to optimize.

At the infrastructure level, there has been a growing trend towards increased consolidation and commoditization of the hardware. Specifically, in contrast to individual silos of computational equipment to handle departmental IT needs, organizations are evolving towards large consolidated data centers. Similarly, advancements such as virtualization and blade servers have further enabled consolidation at the application level.

At the software level, the workload mix is moving away from SPEC-like computing workloads to more generalized, often multi-tier, business processing applications. Similarly, the usage model is moving from predictable batch behavior to more transient spiky behavior. Additionally, with increased consolidation, multiple services are often deployed on the same infrastructure reducing the relevance of physical hardware boundaries. In these cases, instead of optimizing for individual performance of specific workloads, the system design needs to be optimized for utility functions targeted at the service-level agreements (SLAs) of the individual ser-

vices as well as the total cost of ownership of designing, building, and managing the infrastructure.

These changes in enterprise IT architectures have several consequences for system architecture design. For example, increased reliance on mission-critical and online IT infrastructure and greater emphasis on costs leads to a greater focus on issues such as availability, manageability, and scalability. Additionally, greater compaction often leads to greater facilities-level challenges in terms of the need for better thermal provisioning and power consumption. New architectures and solutions are needed that target common-case patterns matched with the emerging mix of workloads and their usage patterns. There also needs to be a greater emphasis on solutions that go beyond simple component or system design to evaluate system-of-systems or utility-level designs. In addition to the need for research on new architectures, there is an equally great need for new models, metrics, benchmarks, and tools in the context of the new class of enterprise workloads and system deployments.

In this paper, we summarize some recent trends in enterprise IT systems and discuss the implications for architecture research. We seek to identify some high-level challenges and open questions for the community to address, and present some data from real-world traces that back the relevance of these approaches.

ENTERPRISE IT TRENDS

This section summarizes some key system and workload trends in current enterprise architectures and some associated implications.

Facilities: Growth of consolidated data centers

Past enterprise IT organizations typically had individual silos of computing infrastructure, each typically targeted at the IT needs of an individual department. Recent trends towards consolidations have led to the emergence of large data centers that house most of the computing and communications facilities associated with the IT needs of the entire enterprise. Such unified data centers provided benefits from real estate consolidation as well as from amortizing the costs needed for facilities security and redundancy, hardware and software consolidation, and lifecycle management and operations maintenance.

A consequence of unifying the IT infrastructure in the same physical location has been an increased emphasis on facilities-level issues. In particular, power delivery, power consumption, and heat extraction are key challenges in the operation of data centers [1]. Though some of these challenges

can be addressed partly by novel methods at the conventional facilities levels [2], they increasingly need intervention at the level of the hardware and software design. Similarly, such infrastructure consolidation has also led to the need for solutions to automate operations management tasks that can otherwise contribute to a large fraction of the total data center costs.

Hardware: Commoditization, modularization, virtualization

As Moore's law provides for increased computational power for the same cost, it has equivalently led to decreased costs for the same computational power. Economies of scale from greater market share and better business processes have also led to further cost reductions. For example, IDC estimates that the overall server average system selling price fell by about 16% from 2001 to 2002 [3].

Another recent development has been the emergence of blade server offerings from most computer vendors. A server blade is essentially a server on a card and represents further miniaturization and modularity from the rack optimized form factor. Blade systems are typically marketed as an enclosure of multiple modular servers with associated benefits of integrated management, control and maintenance.

However, the increased compute density comes with associated problems of power and heat density. For example, it is estimated that future blade servers will increase the power density from current values of about 8KW/rack to almost 55KW/rack causing a subsequent increase in cooling from 27BTU/hr to almost 200BTU/hr [4]. At these levels, this may mean the need for liquid cooling in data centers!

This increased acceptance of blades also has implications on architecture designs in terms of how the individual components (processing, memory, and networking) are organized. For example, many blade designs are also disk-less, indicating another point in the system design space for enterprise IT applications. The increased modularity also opens a new design space of architectures that connect heterogeneous blades in application-specific ways. The increased density with blades also means a greater number of servers in a single data center, with all the associated complexity of management and maintenance of a distributed system with several thousands of nodes.

Finally, server, storage, and network virtualization are becoming commonplace in enterprise IT environments. For example, a recent report [5] summarizes six classes of server virtualization likely to be present in future environments – OS emulation (e.g., Java Virtual Machine), workload management (e.g., job schedulers), distributed workload management tools (e.g., Platform Computing's LSF), hardware emulation layers (e.g., microkernels), resource management (e.g., system partitions), and distributed resource management (e.g., HP's utility data center).

Virtualization leads to greater consolidation and increases system utilization rates. Additionally, the blurring of physical system boundaries through aggressive resource sharing has implications on the choice and design of the specific

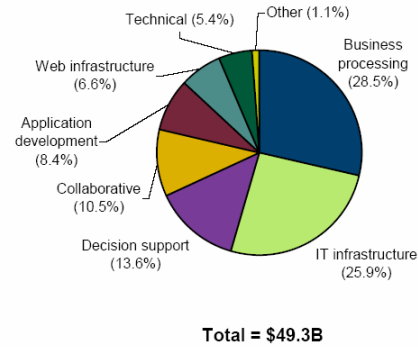


Figure 1: Mix of enterprise workloads classified by revenue [3]

resources. The virtualization layer also provides a layer of control that enables more aggressive system optimizations that were earlier not possible at either the hardware or operating system layers.

Software: Changing workload mix and usage patterns

Figure 1 presents data illustrating the mix of enterprise workloads [3]. Business processing workloads (such as enterprise resource planning, customer relationship management, online transaction processing, and batch jobs) constitute the largest fraction followed by IT infrastructure workloads (file, print, network servers and systems management) and decision support systems (data warehousing and data analysis). In terms of units shipped, IT infrastructure and collaborative workloads (email and workgroup) constitute more than half the servers sold. Scientific/engineering computing is typically a small fraction of the total revenue or units shipped.

While this data is indicative of the high-level trends we discussed earlier – namely the closer tie-ins between IT and business – it also has implications on system design. As is evident from this data, the workload mix is migrating away from specialized technical computing workloads to more generalized, often multi-tier, business processing applications.

Similarly, the workload usage model is also changing. Figure 2 illustrates the CPU resource utilization pattern for two different classes of enterprise workloads. The first application represents a three-tier IT infrastructure while the second one represents a consolidated client infrastructure workload. As can be seen from these pictures, the workload usage behavior is moving from predictable batch behavior to more transient spiky behavior. In all the servers evaluated in this data, the average utilization is significantly lower than the peak utilization. Additionally, looking at the both the workload classes studied in this data, we observe that the sum of the individual peak resource usage per system is significantly higher than the peak of the total resource utilization across the entire solution. These are interesting trends that can be leveraged in future architecture design.

Additionally, when these individual services are consolidated on the same infrastructure, for example in an environment supporting service-centric or utility computing, the

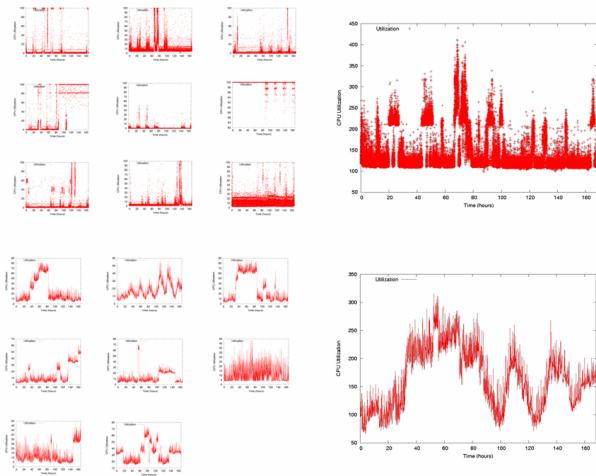


Figure 2: CPU utilization patterns for two enterprise class workloads. The smaller graphs represent the utilization patterns for individual servers while the larger graph represents the utilization pattern for the cumulative solution.

physical hardware boundaries associated with a particular workload tend to get fuzzy and consequently the hardware now needs to be optimized for different combinations of these services. For example, if we were to assume that the two different workloads in Figure 2 were to be deployed on the same infrastructure, the trends discussed earlier – namely the differences between the common case and the peak case and the differences between the sum-of-the-peak and the peak-of-the-sum – are further strengthened.

Metrics: Cost and “utility” over performance

The final key trend in emerging enterprise system design is in terms of the objective function that needs to be optimized. System designs are increasingly being optimized for metrics beyond just performance to also include metrics such as scalability, reliability, availability, manageability, compaction, power, and total cost of ownership (TCO). Typically, these are encapsulated in service-level agreements (SLAs) negotiated with the service that the system tries to achieve. With multiple services deployed on the same infrastructure, the system design needs to be optimized for “utility” objective functions that incorporate the specific SLAs of the individual services as well as the TCO of designing, building, and managing the infrastructure.

IMPLICATIONS FOR ARCHITECTURE RESEARCH

The above-discussed trends motivate corresponding research on new design options for future enterprise architectures. Below, we discuss two key areas: (1) the need for benchmarks, metrics, tools, and models, and (2) the need for architectures to optimize for the key trends in the new class of workloads and systems, with a focus on TCO.

Benchmarks, metrics, simulators, tools, and models

Most of the current research in the community is driven by benchmark suite such as SPEC{int, fp, web}, SPLASH, or more recently, TPC. In addition to not focusing on many of the largest classes of enterprise workloads in Figure 1, these

benchmarks also suffer from the drawback that they focus purely on peak performance. Given the large differences between the peak and the common-case usage, this can lead to significant inefficiencies in design. Thus there is a need for a publicly-available set of consistent traces that track the resource utilization patterns of realistic system deployments that can then be used to optimize future designs.

Similarly, there needs to be greater visibility on the constraints enforced in real-world service-level agreements and utility functions. Irwin et al. discuss a preliminary model to capture utility [6]; however, more work is needed to refine these metrics in the context of architecture design.

Currently new designs are evaluated using performance counters, simulators, or models. However, many of these models do not scale to the emerging trends in enterprise systems running multiple services on systems-of-systems. At the performance measurement level, more work is needed on tools such as Splice [7] that enable aggregated collection of facilities and IT-level data in distributed systems. Similarly, more elaborate simulation models that capture the holistic elements of enterprise system-of-system deployments need to be developed. Finally, more work is needed on models to quantify metrics such as scalability, availability, security, manageability, etc.

Architectures optimized for the “new common case”

A lot of architecture research is predicated on optimizing for the common case. We expect this to continue, albeit in the context of a better understanding of what the common cases are in emerging enterprise scenarios. In particular, we believe that future enterprise architectures will leverage the following common-case trends.

Optimizing for the nominal versus the peak: Understanding the differences between the peak and nominal (common-case) behaviors in terms of resource usage can have significant implications on the architecture design. As a simple illustration, designing a power supply that has its maximum operating efficiency at normal load instead of at the peak load can often yield significant total energy savings. Architectures that are designed to match the inflection points of the technology efficiency curve with the common-case behavior of the workload are likely to have the best cost-utility tradeoffs. An example of this approach is the idea of heterogeneous multi-core architectures [8].

Optimizing for system-of-system versus individual system: As seen in Figure 2, there is often a significant difference between the behavior of individual systems and that of the total system-of-systems. As enterprise architectures evolve to multi-tier, multi-service, consolidated solutions, architectures that optimize for an objective function targeting the entire system-of-systems should be more efficient than architectures focusing on individual systems (e.g., [9]).

Multi-level power management: Several of the trends discussed earlier – facilities consolidation in data center, increasing hardware modularity, and greater compaction – all motivate the need for power management as a first-class

design constraint for enterprise architecture design. Though several promising studies have evaluated power management for server architectures [10,11,12,13], more work is needed for the increasingly important power constraints of the future.

Diversity-aware architecture designs: One of the key recurring patterns in our earlier discussion of enterprise trends is the potential diversity in future systems – at the level of the workload, at the level of the service, at the level of the resource demand requirements, and at the level of the SLA constraints. Future enterprise architectures will need to be designed to perform better under such diversity. Simple, modular, general-purpose designs that provide hooks for adaptivity will likely perform better under these constraints. A key challenge that will need to be addressed is the increased verification complexity with such added diversity.

Architectures optimized for reliability, manageability, security, and availability: Optimizing for these metrics can have a wide range of architectural implications. Diva [14] is an example of reliability enhancement at a low level. Intel's announcement of the Vanderpool and Silverveil support for virtualization indicate some other potential architectural enablers for these features.

Commoditized building blocks: Another key area for architectural innovation is to leverage commoditization and scaling for lower cost solutions. This may argue for higher-level solutions based on multiple simpler architectures. FAB [15] is an example of building reliable and scalable storage systems out of commodity building blocks.

Finally, optimizing any one of these common cases without regard to the others is insufficient. Instead, all of the above must be considered within a framework of minimizing TCO.

CONCLUSIONS

In summary, enterprise IT systems have seen many significant changes in recent years. At the facilities level, large data centers have supplanted individual departmental-level compute clusters. At the hardware level, trends favoring commoditization, virtualization, and modularity have led to new challenges with increased resource consolidation. At the software level, multi-tier workloads with spiky usage patterns have displaced conventional batch-based predictable applications. As systems evolve towards service-centric and utility computing environments, there are likely to be new challenges with increased workload diversity and resource sharing. Along with all these changes, there has also been a gradual shift in the objective function being optimized – from a pure performance focus to a broader SLA-driven TCO focus.

These dramatic changes motivate corresponding changes in the emphasis of architecture research targeted at future enterprise systems – both at the level of the benchmarks, metrics, tools, and models to facilitate such research, as well as the architectural level that designs are optimized for these new constraints. In this paper, we discussed a few of the salient trends and potential architectural implications. While

fundamental architectural principles like simplicity and optimization for the common-case are likely to continue, we believe that an intelligent re-evaluation of these principles in the context of the new systems is necessary. The designs articulated in this paper, such as architectures optimized for TCO, architectures optimized for nominal behavior of system-of-systems, or architectures optimized for diversity only scratch the surface of what is possible. We believe that these areas offer a rich opportunity for more innovation for the broader architecture community.

REFERENCES

1. W. Tschudi et al., Data Center Energy Research Roadmap, Lawrence Berkeley National Laboratory Report, 2003.
2. C. D. Patel, A Vision of Energy-aware Computing from Chips to Data Centers, Keynote address at the Japan Society of Mechanical Engineers, International Symposium of Micro-Mechanical Engineering (ISMMME), 2003.
3. M. Eastwood and J. Yang, Worldwide Server Workloads 2003: The continuously evolving relationship between applications and servers, IDC #30686, Volume 1.
4. J. Mouton, Enabling the vision: Leading the architecture of the future. Keynote speech, Server Blade Summit, 2004.
5. T. Bittman, The Future of Server Virtualization, Gartner Research Note, T-20-4339, July 17, 2003.
6. D. Irwin et al, Balancing Risk and Reward in a Market-based Task Service, Thirteenth IEEE Symposium on High Performance Distributed Computing, June 2004.
7. J. Moore et al, A Sense of Place: Towards a Location-aware Information Plane for Data Centers, Hewlett Packard Technical Report TR2004-27.
8. R. Kumar et al, Single-ISA Heterogeneous Multi-Core Architectures: The Potential for Processor Power Reduction. In Proceedings of the 36th Annual International Conference on Microarchitecture, December 2003.
9. D. Irwin et al, "Dense and powerless: hardware-software coordination for future blade designs," HPL Technical Report, HPL-TR-2004-33, 2004.
10. J. Chase et al, Managing Energy and Server Resources in Hosting Centers, In 18th Symposium on Operating Systems Principles (SOSP), October 2001.
11. M. Elnozahy et al. Energy-efficient Server Clusters. In Proceedings of the Second Workshop on Power-aware Computing Systems (PACS), Feb 2002.
12. E. Pinheiro et al, Load Balancing and Unbalancing for Power and Performance in Cluster-based Systems, In Proceedings of the Workshop on Compilers and Operating Systems for Low Power, September 2001.
13. J. Moore et al, Going beyond CPUs: The Potential for Temperature-aware Data Centers, Workshop on temperature-aware computing systems, 2004.
14. T. Austin, "DIVA: a reliable substrate for deep submicron design," In Proceedings of the 32nd Annual International Conference on Microarchitecture, pages 196-207, 1999.
15. Y. Saito et al., "FAB: building reliable enterprise storage systems on the cheap," In Proceedings of ASPLOS XI, 2004.