

White Paper



Maximizing returns on IT investment with end-to-end performance management

Rajinder K. Gandotra

Organizations, in the quest to enhance operational efficiency, develop additional functionality through customizations and new product deployments. This results in development/deployment of application silos and, at times, additions to hardware. Such additions lead to unmanageable application/infrastructure with little focus on performance. Also, CIOs and CTOs are under constant pressure to deliver more with less investment. This article proposes an end-to-end focus on performance management that can maximize returns on existing investment by optimizing IT infrastructure

Organizations are under constant pressure to deliver better value to their customers and stakeholders. They typically respond by developing additional functionality through customizations and new product deployments. This results in development/deployment of application silos and, at times, additions to hardware. Such additions lead to unmanageable application/infrastructure with little focus on performance. The problem is compounded by the tremendous increase in size and complexity of applications and infrastructure. Further, in the current economic scenario, CIOs and CTOs need to reduce the TCO of all systems and deliver more value with less investment. This paper discusses how organizations can maximize returns on new and existing investment by optimizing IT infrastructure through end-to-end performance tuning and management.

Functionality-centric focus, combined with paucity of resources and time, results in faulty design of application or coding style. In addition, application testing is often conducted in a scaled-down environment to test only functionality, not performance. This results in performance problems under high workload conditions in production environments. Firms may also be disadvantaged by damaged customer relations, lost income, increased hardware costs, increased development costs, and cancelled projects.

An end-to-end performance paradigm with special attention to performance and improvements to existing deployments can be effectively used to address performance issues. Benefits from such an approach are manifold and include high ROI.

End-to-end performance paradigm

The end-to-end performance paradigm focuses on performance at every stage — planning, design and deployment; and optimizes all layers of infrastructure and applications on an ongoing basis. This approach provides a holistic view of the system from planning and deployment to usage and constantly evolves it to meet dynamic customer expectations/performance targets (Figure 1).

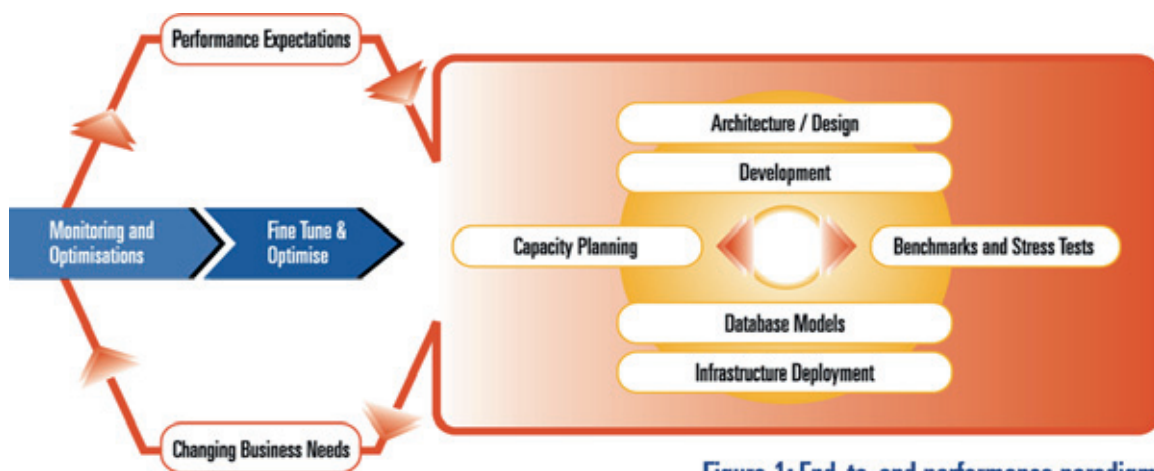


Figure 1: End-to-end performance paradigm

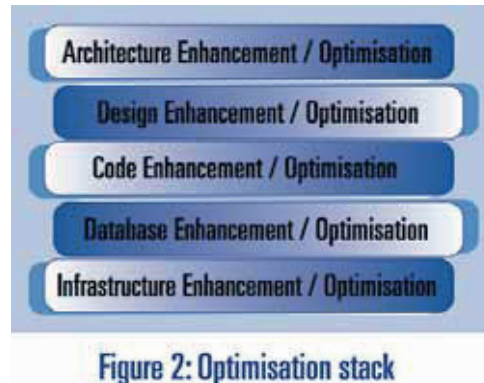
The end-to-end performance paradigm requires organizations to:

- Set appropriate, reasonable, tangible, measurable and justifiable performance goals

The solution to performance problems lies in the clear understanding of stakeholder requirements and user perceptions. Establishing a tangible, measurable and realistic performance goal is of utmost importance as all other aspects evolve from it. Obviously, the performance goal must be set in advance to address performance issues effectively.

- [Architect, design and develop from a performance perspective](#)

Figure 2 depicts the five layers in performance optimization. Performance focus at the architecture and design layers are the most rewarding because re-working these layers after system development and deployment is timeconsuming and expensive. Since customer needs and technology are dynamic in the competitive market, regular performance reviews of the design and architecture are necessary.



- [Allocate adequate capacity based on benchmarks / stress tests](#)

This can be achieved by:

- Adequate capacity provisioning
 - Designing and deploying IT infrastructure for optimum performance
 - Capturing changing business requirements
 - Monitoring, fine-tuning and optimizing on an ongoing basis
 - Being ready for the future
- [Design and deploy IT infrastructure for optimum performance](#)

It is important to implement learnings from benchmarks/stress tests in the final deployment. Infrastructure layers such as servers (CPUs, memory and I/O access bus, cache in relation to application), storage architecture (disk drive type and speed, cache on storage), RAID type, I/O access mechanism and parallelization, network (bandwidth and access between servers), and WAN should be evaluated holistically so they do not become performance bottlenecks.

- [Capture changing business requirements](#)

Each system is designed to serve a set of well-defined expectations. Each application should therefore be used in accordance with its intended use. The impact on the overall system must be analyzed before making changes to applications.

- [Monitor, fine-tune and optimize on an ongoing basis](#)

Proactive and ongoing monitoring is important to meet ongoing performance requirements. Regular performance audits and tuning improves performance and availability.

Performance tuning is a highly skilled task, requiring expertise and experience in all the five layers discussed earlier. A holistic understanding of the system helps deliver best performance tuning results. Domain knowledge is vital for quick and effective performance tuning. With increase in complexity, specialized technical resources and services must be involved to help the organization address performance issues effectively. The most effective way to optimally tune the system is to have experts review system performance regularly. Regular tuning exercises can deliver improvements up to 60% (initial tuning attempts show higher improvements).

Be ready for the future

Regular trend analysis and projections on system usage patterns and changing business needs can determine additional resource requirements. This avoids surprises and the need for ad hoc purchases and siloed deployments. It also provides sufficient time to negotiate deals, as there is no pressure to upgrade or deploy at short notice. Leveraging performance from a deployed system is not a destination but a journey. Typically, deployments evolve in terms of usage, process and complexities. Market forces change business requirements, which in turn change system user expectations. These changes alter the importance of business functions or result in new functionalities to meet business requirements. For instance, user transactions with banks have changed dramatically in the last few years, from branch banking to on-the-spot delivery channels.

Fierce competition has made optimally performing systems mandatory. This is evident from the increasing focus on IT governance, SLA management, and Quality Of Service (QOS) on performance aspects. An optimally performing system increases customer loyalty by increasing overall availability and satisfactory and reliable response from the system. It also reduces overall investment on infrastructure by optimally utilizing existing resources. The end-to-end performance paradigm can be effectively used to leverage high performance and scalability from deployed IT solutions and ensure higher returns on existing investment.

About the Author

Rajinder K. Gandotra heads the Performance Engineering and Enhancement Practice at Infosys. He has over 15 years of experience in the IT industry.



For more information, contact askus@infosys.com

About Infosys

Many of the world's most successful organizations rely on Infosys to deliver measurable business value. Infosys provides business consulting, technology, engineering and outsourcing services to help clients in over 30 countries build tomorrow's enterprise.

For more information about Infosys (NASDAQ:INFY), visit www.infosys.com.