To be presented at CMG 2000, December 10-15, 2000

eBusiness Performance: Risk Mitigation in Zero Time

(Do it Right the First Time)

Amy Spellmann
HyPerformix, Inc.
amy@hyperformix.com

Richard L. Gimarc
HyPerformix, Inc.
rgimarc@hyperformix.com

As eBusiness grows exponentially, Web performance risks will skyrocket. Corporations are pressured by competition and consumer demand to produce new applications rapidly. Increasing complexity threatens on-time delivery for organizations that defer performance analysis to the end of the development lifecycle. With potentially conflicting goals, how do business owners and the IT organization work together to overcome the exposures caused by performance problems? This paper combines technological and cultural processes to meet performance goals without impacting the development schedule.

1 Introduction

Most organizations have teams and tools to manage performance after applications are deployed. These are remnants of the client/server era where point-solutions were readily available. Client/server applications that spanned one or two servers were well suited for point performance solutions (e.g., buy another server). However, those days are long gone. The acceptance of the Web as a primary technology of doing business has created a new class of multi-tier systems with complex interactions that challenge traditional analysis. Further, the end user is no longer just the corporate staff, but the customers who have a wide variety of choices in the Web's electronic market place.

Companies that deliver Web-based applications face time-critical challenges:

- Implementing new application features rapidly to beat the competition.
- Maintaining customer satisfaction and loyalty.

The paramount need is to develop and deliver new applications quickly that meet business-defined performance criteria. The risks are high (and expensive) when good performance engineering practices are not incorporated into the application development lifecycle. However, including those practices often implies that projects will take longer to complete, milestones will be missed, and schedules slipped. Although delay is unacceptable to the business owners, so is poor performance. So the notion of “mitigating risk” (especially performance problems) with established performance engineering practices must be sold to and adopted by organizations as a process that adds “zero time” to the development lifecycle.

It must be emphasized that the concept of “***zero time***” relates to the end-to-end system lifecycle, not to the total development resources required to build and test a system or to the time it takes to build a version of the system that does not perform adequately. The timer starts when the application is specified by the business owner and ends when the deployed system’s performance meets the business owner’s performance requirements.

As an example, consider Figure 1, an illustration of the software lifecycle. Project A was developed without performance engineering. After deployment the application architecture had to be reworked twice before performance was satisfactory (this is a real-

Project A	Development Lifecycle	Rework 1	Rework 2
Project B	Development Lifecycle (with Performance Engineering)		

Figure 1. Project Duration with and without Performance Engineering

world example). Project B was developed using performance engineering practices with additional dedicated performance engineering resources. At deployment, system performance met the requirements specified by the business owners; performance engineering added “zero time” to Project B.

Case in point, an application development group was charged with implementing a Web-based training system for teachers: distance learning. They were concerned about performance and wanted to incorporate architectural analysis, load testing and predictive modeling into their software development lifecycle. When asked about the development schedule, they said that they had exactly two months to complete the project – from start to finish. "How in the world," they asked, "can we follow any good performance engineering practices in two months?" They preferred to re-architect and fix problems later rather than slip the schedule.

Slipping the schedule is not an option for 99% of Information Technology (IT) organizations; business owners are demanding new applications at a torrid pace. As a result, corporate culture rewards development organizations for fast delivery; and often rewards them a second time six months later for fixing the very performance problems they overlooked during development. No wonder development organizations do not want to change their processes. Clearly the focus has to be on cultural change and institutionalization of performance engineering. All the involved parties must be convinced that this engineering process assures performance and does not lengthen the development lifecycle.

This paper describes the cultural changes required to meet performance goals in zero time and the supporting technologies needed at each step of the system lifecycle. The focus is on the success of the business, the performance certification of the deliverables, and the integration of performance engineering methodologies and tools.

2 The Case for eBusiness Performance Engineering

Performance engineering is not a new science; there have been numerous publications on the subject, like Connie Smith’s excellent treatise on software performance engineering [SMIT1990]. IT organizations have been slow to adopt these practices because of the assumption that performance engineering increases the time required to deliver software systems. In the Internet world, performance engineering has become critical for eBusiness. In this paper, eBusiness refers to networked systems that interact directly with “customers” to provide sales and services [BART2000]. The term “customer” refers to both internal and external customers. An eBusiness strategy is complex and pervasive, encompassing both internal and outward-facing systems. Both external customers and internal processes require good performance for “customer” satisfaction and staff productivity. Business owners and IT must work together to deliver on these strategies.

Why should business owners, architects, developers, quality assurance (QA), and operations agree to the institutionalization of performance engineering? The answers have to be clear and compelling for each organization. To characterize each message, first look at the groups interrelationships and goals. Figure 2 illustrates the difference in the focus and goals of these groups.

- The goal of **business owners** is to run a profitable business. These owners look to the IT organization to provide the systems needed

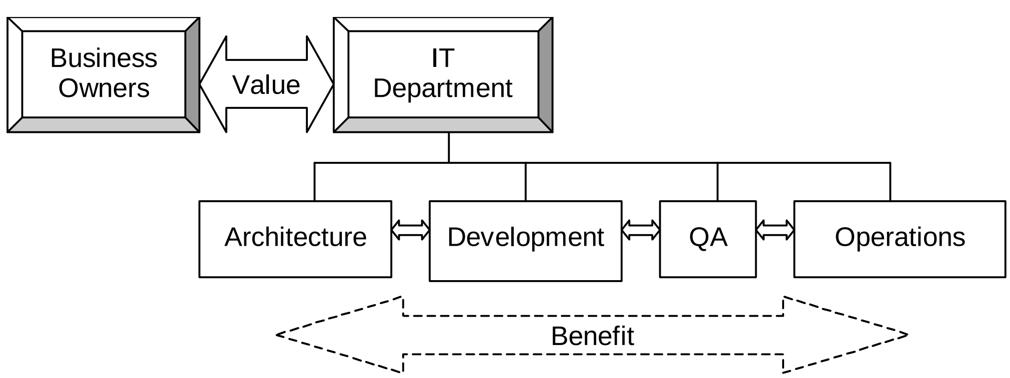

Figure 2. Benefit versus Value

to deliver services to the customer. They are looking for **value**, estimated worth, market price. Business owners must work with IT to determine the return on investment (ROI) for dollars spent in developing and maintaining the systems that provide competitive applications.

- The goal of the **IT department** is to provide a well-engineered system that is cost-productive to build and to maintain. For the IT groups, the **benefit** of performance engineering counts most, enabling them to meet their goals.

The key to making performance engineering work is to address the value and benefit issues of each organization. Business owners get more value for their dollars if the system performs on deployment rather than having to face customer wrath and desertion caused by catastrophes that inevitably occur. There are numerous examples of catastrophes in the market place: eToys' Christmas fiasco, eBay's auction flop, and E*Trade's inability to meet customer demands during market storms. To avoid these problems, there must be synergy between the business owners and IT; but IT marches to a different drummer and does not always align itself with the demands of the business. Better communication and education are essential to improving the value provided by IT.

IT executives have a mandate to manage complexity and allocate scarce resources only to essential activities. They do this by promoting a bias for results and developing organizational excellence. IT executives must create convergent team thinking, and challenge and change the culture of their organizations. In IT organizations today, architecture, development, QA and operations have a loose relationship to one another. The architects develop a design and "throw it over the wall" to development; development does a similar handoff of code to QA; and QA hands off the system to operations. In order for performance engineering to benefit the entire IT organization, there must be a protocol to handle intra-organizational communication. Performance certification provides a means for demonstrating and communicating compliance with performance requirements.

3 Performance Certification

Business owners and the IT organization must agree that performance is key to the rapid deployment of new applications.

- Business owners will agree if they can be given assurance that the deployed system will meet their performance requirements.
- IT will agree if the quality of their delivered product and the effectiveness of their staff improves.

In both cases, it must be demonstrated that performance engineering is a complimentary discipline that reduces risk, increases confidence, and supports aggressive delivery schedules.

Estimating the dollar value of performance can be difficult and varies widely based on business application, value of the application to the business, cost of development and deployment, etc. One example illustrates this clearly; an Intranet procurement system being developed by a large bank. This system is expected to save the bank $20M per year. Cost to implement is approximately $3M. The cost of failure is that the bank will not be able to reap the $20M per year savings and will have already spent $3M for implementation.

3.1 System Lifecycle

The business owners begin by describing the requirements of the new application. These requirements are passed to the architecture group that creates a design to support the requirements. The development group then writes the code according to the design and passes it on to the QA group which tests it and determines when it is ready for production. The operations group is then challenged with administering and managing the system.

How do you assure that through this process the system satisfies the requirements of the business owners? First, the business owners must define performance requirements. These requirements must be consciously carried through the lifecycle. At each handoff between organizations, it must be demonstrated that the deliverables satisfy the stated performance requirements. Performance issues must be identified, analyzed, understood, and addressed step-wise through the lifecycle. The benefit and value of this process is that the delivered system performs as required with supporting tools and methodologies in place for continuing operational success.

How do you add performance awareness to the lifecycle to:

- Demonstrate that performance requirements are satisfied at each step.
- Increase the confidence of downstream organizations.
- Benefit each organization.
- Stay on schedule.

Performance certification provides a means for communicating compliance with performance requirements. Figure 3 shows how performance certification (top arrows) integrates with the system lifecycle deliverables (bottom arrows). Performance requirements are passed down the chain with the deliverables and each organization knows precisely what they have to do to meet those requirements. Business owners are at the head and tail of the lifecycle since they develop system requirements (functional and performance) and are also responsible and accountable for the success or failure of the deployed system.

The deliverables at each step are backed by a "performance certification". Instead of simply throwing deliverables over the wall for the next step, each organization must demonstrate that their deliverables support the stated business-level performance requirements. The deliverables must be defensible from a performance perspective

3.2 Performance Certification Resources

Performance certification requires a breadth of involvement that is not normally found in IT organizations. Most organizations have pockets of performance skills, but these are generally isolated in nature. For example, it is common to find the capacity planning group within operations. This group generally has a long-term and coarse view of the world: they are viewing the system as a whole, not always focused on individual applications. There may also be separate performance analysis groups that spend the majority of their time fighting "performance fires".

As described below, performance certification requires a steadfast awareness of performance issues throughout the system lifecycle. It is unreasonable to place this extra burden on the individual IT organizations; they have neither the skills nor the time. Instead, we advocate the creation of a separate Performance Engineering group chartered with the task of performance certification. The two case studies presented in Section 5 describe such an organization.

3.3 Business Owner Certification

In addition to stating functional requirements, the business owners must also describe application-level performance requirements. For example, the business owners at an online brokerage firm might set response time goals for various arrival rates:

- 2 second response time per trade for an arrival rate of 100 trades per minute.
- 5 second response time per trade for an arrival rate of 2,000 trades per minute

These requirements not only provide guidelines for the operational behavior of the system, but also give the business owner the means to quantify the value of the system in terms of dollars per minute. To develop these requirements, the business owners must forecast expected end-user load on the system and predict growth of the user base. Analysis of business metrics and growth trends will be necessary to accomplish this.

3.4 Architecture Certification

The goal of the architecture group is to create a system design that satisfies the functional requirements laid down by the business owners. The new performance perspective further requires that the architecture group provide a high-level architectural model to demonstrate that their design meets performance requirements. In addition to

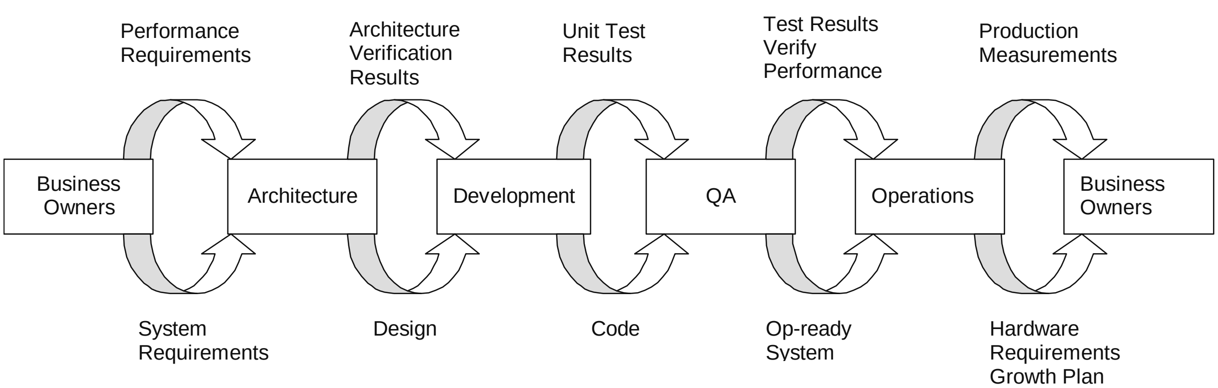

Figure 3. Performance Certification Program

demonstrating system-level performance compliance, the model can be used to determine component-level response time constraints and the metrics to be measured and reported during development and in production. Simple stand-alone server analysis can be performed using analytic techniques or a spreadsheet. As system complexity increases to a multi-server configuration, simulation modeling has proven useful [NEUS1998]. During the architecture phase, a high-level model is sufficient for evaluating design choices. The architecture group can then prove the feasibility of their design.

A large consulting firm serves as a good example of architecting a new system. Their implementation of a knowledge base (a central repository for all of their products, lessons learned, and customer information) was considered to be a performance-critical system. This system had to be architected to perform well on pilot release – early perception of performance problems would set a bad precedent and impede the project's success. Instead of waiting until development was complete, they used a high-level model of their multi-tiered Intranet application to estimate performance. This gave them a method for demonstrating performance feasibility of the system to the developers as well as their sponsors.

Additional benefits to the architecture group include:

- Confidence that their design satisfies business owner requirements.
- Ability to provide feedback to the business owners if performance goals are not achievable.
- Insight into the performance requirements of application components.
- Basic understanding of the sensitivity of architectural components.

3.5 Development Certification

Performance certification helps the development organization focus their efforts. Development is tasked with implementing the system. In addition to the functional behavior, the developers must demonstrate that their implementation meets the performance criteria stated by the architecture group. Before development begins, they will know that the business owner's performance requirements can be satisfied due to the results/certification provided by the architecture group. They will also have component-level targets to monitor during development. This gives them guidance on instrumentation as well as goals for unit performance tests. Measurement results from unit tests will prove that their code meets the architecture group's component-level performance guidelines.

Additional benefits to the development group include:

- A methodology for basic measurement and reporting to be used to demonstrate compliance with the architecture group's component level performance objectives.
- A guarantee that the basic component-level performance requirements are met.

Performance certification does not burden the development staff with the addition of measurement and reporting. It focuses their performance tuning efforts on the sensitive components identified by the architectural analysis. It requires that they deliver a system that meets both functional and performance specifications.

3.6 QA Certification

The QA group is responsible for testing the functionality and performance of the system they receive from development. This testing can be done synergistically to achieve maximum benefit and have minimum impact on schedules.

From development, QA has been given an application and a basic set of application-level measurement tools. Their job is to perform large-scale load testing and/or modeling to verify system performance under load. This is the final chance to verify that the delivered system meets the stated business requirements before going into production. Frequently, models developed by the architecture group are used as a starting point for the more detailed models required by QA. These models can be used to increase the scope of load tests and help focus their effort. For example, it may be physically impractical to test the new system under conditions of extreme user load. This scenario can easily be evaluated using a model of the system. The level of system complexity guides the amount of modeling and load testing required. For a single-server system, a simple load test will often be sufficient to demonstrate performance. While for multi-tiered systems, a combination of modeling and load testing may be necessary to increase coverage. Again, the amount of effort required can be reduced by reuse of models used by other groups in the lifecycle chain as well as combining functional and load testing automation.

The benefits of performance certification to the QA group include:

- Proof that the system they hand off to operations will meet the business owner's requirements.
- Validation of the measurement methodology defined by the architects and implemented by the developers.
- Coverage of extreme load scenarios to include those that cannot easily be tested.

3.7 Operations Certification

When the operations group receives the complete system, they will have the assurance that the system performs as specified. Operations can then focus on doing their job: installing the software, integrating the new application with existing systems, planning for growth, and implementing new reporting methods. These tasks insure that the operations group can report back to the business owners that infact, the system does perform as required.

Operations will be immediately ready to monitor and report on the customer experience. Additionally, they can leverage the tools and methodologies built by the other IT groups to help them to plan for growth and continue meeting performance goals. Operations may have to revise and update these tools to account for changes in the application and system environment. Additionally, operations may incorporate new monitoring tools along with those provided by development and QA. All of these tools support operation's ability to manage the IT systems more efficiently and effectively.

Performance certification benefits to the operations group include:

- Reduced firefighting.
- Improved process for monitoring the customer experience.
- Improved tools and techniques for capacity planning.
- Proven methods for demonstrating compliance with the business owner's performance requirements.

3.8 Limited Certification

It is important to note that performance certification need not be performed at every step in the development lifecycle. There are ways to streamline in all cases. For example, in the architecture phase, a model may be developed in a day – very simple representations can be extremely valuable for understanding the dynamics of a system. The appropriate tool and time requirements must be understood in each case. Spending 6 months modeling an Internet trading application that has to be delivered in two months makes no sense. Nevertheless, some sort of analysis must be done to verify that the design meets performance goals. The level of analysis performed will be based on complexity, business goals, and resource availability. The certification process must not adversely effect the end-to-end time to deliver the new system.

4 Institutionalizing Performance Engineering: The Performance Plan

A performance plan results from the integration of the certification process with the system lifecycle. The "plan" is the road map for tracking performance requirements through the development process. In addition, the tools and techniques utilized to demonstrate performance compliance are precisely those that you need in place to successfully deploy and manage the new application. The effectiveness of performance engineering can be demonstrated by using a performance plan to focus performance awareness through the system lifecycle. The ultimate goal is to institutionalize performance engineering in all projects and system development activities. This is accomplished through a combination of the integration/reuse of technologies and education.

On the technological side, institutionalization requires the integration of data, tools, and methodologies to derive maximum benefit from the technology and the process. It is imperative that the process not be a burden. High-level models developed during the architecture phase can be refined for use later in the lifecycle (QA load testing and operations capacity planning). Application instrumentation adds value at all steps in the lifecycle by providing application-level measurements that assist development, QA, and operations. The combination and refinement of measurement and analysis tools and methodologies provides the IT staff and business owners with precisely the set of information they need to monitor application performance today, and plan for the future.

Integration of data, tools, and methodologies results in a more streamlined approach that reduces the time required to perform the certification. Thus the development schedule shrinks even further. Figure 4 depicts the relationships between tools, modeling and measuring, and their interaction throughout the project lifecycle. Incorporating data from development's unit tests can refine models developed during the architecture phase. These models become part of the QA process with some

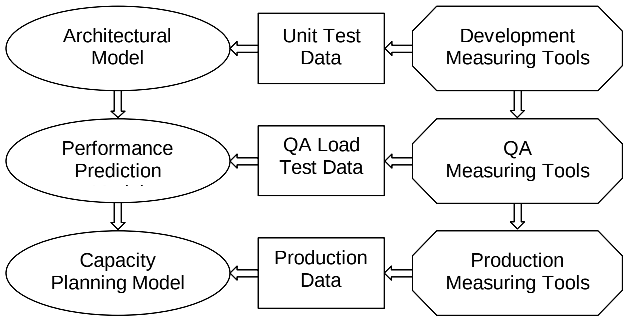

Figure 4. Interaction Between Tools, Modeling, and Measuring

refinement and the addition of data from QA load testing. The result is a capacity planning model that the operations group can use for evaluating future system performance under different loads and configurations. The "tools" used for performance certification are refined at each step in the lifecycle. At deployment you have a well-stocked arsenal of application-specific performance measurement and analysis tools.

Along with the technological integration of tools and data, education and communication are essential to institutionalizing performance engineering. The demonstrated value and benefit must now be communicated throughout the company to demonstrate to other business owners and development organizations that this process improves everyone's life. Performance awareness classes should be offered to management, business owners, architects, development, QA, and operations. Additionally, specific performance engineering classes should be used to develop the skills needed at each step in the development lifecycle.

5 The Performance Engineering Practice

As mentioned earlier in this paper, the practicality of providing performance certification requires a dedicated performance engineering organization. The primary goal of this new organization is to provide a level of assurance that applications, systems, and services provided by IT satisfy the performance requirements of the business owners. Thus, performance engineering can be viewed as a process to mitigate risk. ***Innovation and change are required to reshape established IT processes and organizations into ones that are more dutifully aligned with business goals. And, IT must be able to view those same goals in terms of the systems they build.***

As pointed out by Lewis [LEWI2000], innovation has three main ingredients: technology, funding, and an entrepreneur. These are exactly the same ingredients needed for a successful performance engineering practice. The remainder of this section contains two case studies describing the creation and growth of performance engineering practices.

Two large companies are currently putting together performance engineering practices, groups that provide performance engineering services to their respective companies. These groups have been formed due to the complete buy-in and sponsorship of top IT executives. Two steps are required for a successful performance engineering practice. The first step is to acknowledge the importance of performance to the business and the value of assuring that systems will be better able to meet performance requirements with the use of performance engineering services. The second step is to institutionalize performance engineering as a formal practice. For business owners, the key is the return on investment. Already, these Performance Engineering (PE) organizations are making a big difference in the quality of applications and systems that support each business.

5.1 The Insurance Company's PE Practice

The first business, an insurance company, started its PE group in early 1997. The group was sponsored by the CIO and other IT executives. Their charter is to provide proactive performance engineering services for the company. Since their inception, the group has grown from 3 performance engineers to a staff of 15. The organization works with other groups inside the company to help them ensure that new applications are designed, developed and tested to meet performance goals. Performance Engineering has come into focus for the IT organization especially now that the company's clients have direct access to their infrastructure through Web-based applications. Performance problems not only effect the productivity of the insurance company's staff, they are apparent to their customers.

Institutionalizing performance engineering has occurred with much dedicated effort to evangelize the power of performance engineering. The company required an internal marketing and education plan to develop support and to build trust from other internal groups. Formal seminars, informal lunch meetings ("brown bags") and presentations at senior management meetings are done on a regular basis. Educating executives and project managers on the value of performance engineering is crucial to the success of the organization. The group has to sell its own services to thrive and survive.

When the insurance company's PE group was first formed, their focus was primarily on performance modeling and testing after the system was developed. This is typical of most companies; their performance concerns typically start after the application is released or about to be released (the time of greatest fear). This required coordination with various testing groups to get data required for modeling the applications and to run load testing scenarios. Once performance project goals were achieved, even at this late stage of the system development, the performance engineers publicized their services and successes. This proved extremely effective in garnering attention from managers of other application development efforts. By early 1999, the PE group was getting more requests than they could handle. But PE group managers were also beginning to have some problems delivering answers fast enough to make a difference. ***When the answers come just before the system is to go into production, it's really too late.*** Other challenges included limited resources from the development organizations; the people on the critical path were always the ones who knew most about the application. PE requires support from these other organizations to get the information necessary to analyze, model and test the system.

As the publicity and advertising continued, the PE group pushed hard to get involved earlier in the system lifecycle. They convinced upper management, application development, and architecture that this is doable and adds more value to their services. Getting answers before the system is developed is key to the concept of *zero time*. The services provided were still basically "probes" on the lifecycle path. There was not a continuous performance certification process in place. With additional successes, especially on highly visible Web applications, they have widened their focus to providing performance analysis services much earlier in the design and continuing up to the deployment stage.

Today, when a new project is started, the project manager works with the PE group to determine the level of risk expected for the new product. If there is a high risk factor, as for most Internet applications, the PE group works with the project manager to gather information that will determine what parts of the performance certification program are appropriate for the system, i.e., fit the schedule, the requirements, etc. In high-risk projects, the PE group starts by working with the business owners to determine performance requirements, volumes and usage patterns for the system. Next, they work with the architecture group to ensure that design decisions for these new applications are made to support the performance requirements. This typically involves a modeling effort, often spreadsheets and discrete event simulation. PE's early involvement ensures that the selected architecture will perform, taking pressure off the development groups to "tune" an architecture that is not designed with performance in mind. The development organization writes the code according to the architecture group's design and may add instrumentation, but often do not. This is an area where the company would like to see better performance engineering practices. A big factor is the lack of vendor-provided tools for instrumentation.

Once the code is written, the PE group is given access to test environments and key application experts in order to test and perhaps do additional modeling for the system. This is the end of the involvement for the PE group at this company. Once load testing and modeling are completed in the QA stage, the system is handed off to operations. Operations does not reuse any of the tools built by the PE group. They would like to get an integrated performance plan that includes capacity planning, but typically they do not work together.

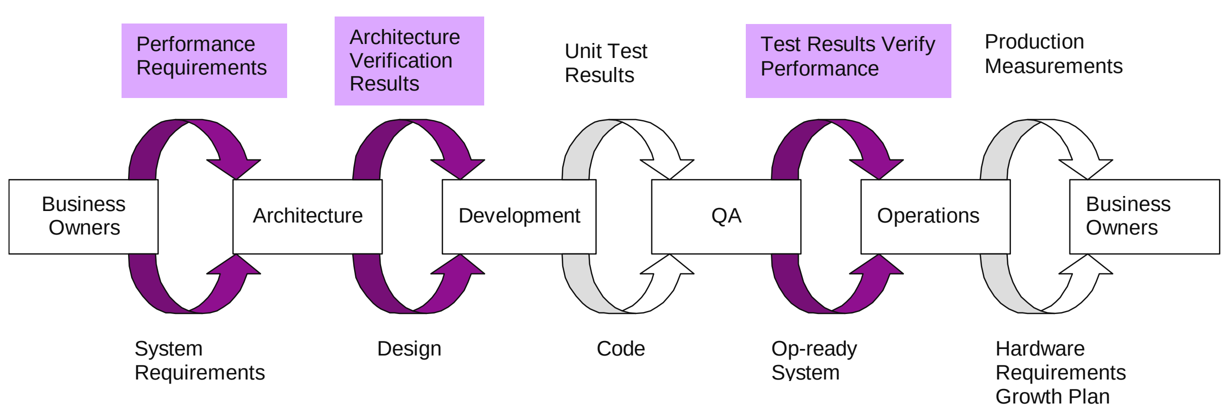

Figure 5. Performance Certification Program Example

In Figure 5, the darker shaded arrows and boxes illustrate the portions of the performance certification program that the insurance company has implemented. PE is working to apply performance certification coverage to the entire lifecycle, including current work on a model framework for their entire IT infrastructure. The model would provide the ability to rapidly assess the performance of new applications in their Internet systems. It would be part of a performance certification plan that includes performance engineering at each stage of the lifecycle. This would include reusable tools for data collection and analysis and framework models of the complete network, with detailed models of specific applications.

A major new system that the company has been working on, a new online banking application, is the clear beneficiary of the PE groups' hard work. The initial step of characterizing performance requirements with the business owners brought performance to the attention of all groups involved in the development. Performance modeling in the architecture and design phase demonstrated that a few changes to the architecture would better position the application to meet performance goals. The PE team, demonstrating their ability to help the project stay on track, performed all of these tasks quite rapidly. Load testing conducted on the online banking application affirmed that the application performance was on track.

The insurance company's PE group has proved that analyzing performance early in the lifecycle ensures that systems are built right the first time. In the past they have had many systems that had to be re-architected and rewritten more than once. Using their proven analysis, modeling and load testing techniques, the group helps to ensure that the systems, including their online banking application, are built right the first time.

5.2 The Financial Services Company's PE Practice

The second company, a large financial services business, began its Performance Engineering practice in April 2000. The company's business relies heavily on Web-based access to their financial services. Market storms, new services, and increased volumes stress the infrastructure and threaten the company's ability to provide their customers with good, consistent response times. Fire-fighting groups cannot proactively work on performance engineering; they are overwhelmed with day-to-day problems. Fire-fighting is past the point of being able to keep up. A proactive approach to performance problems has come to be essential to the success of the business.

This PE group is on a much faster track than the insurance company: they are under pressure to deliver value quickly. The financial group's charter is "to support design for performance excellence throughout the lifecycle of [the company's] Technology projects." In April, the Executive VP of System Technology chartered an autonomous PE group to provide services for company organizations building and maintaining IT systems. As of September 2000, the financial group has 8 members and is building to 10. Their manager believes he needs 20 to support all of the ongoing projects at the company. Already they are getting more requests for services than they can handle.

The financial company's PE group will provide the following services:

- Identify performance constraints and guidelines with business owners.
- Develop performance plans for certification in the architecture stage.
- Estimate performance during development.
- Review test plans and test results.
- Certify performance at deployment.
- Resolve performance problems in production.

The PE group at this financial services company is creating partnerships with all of the diverse groups in IT that have similar charters, e.g., capacity planning and performance testing. They are also developing partnerships with business owners. The initial focus is on their mainframe systems that support Web-based financial services, but is expanding to include all of the online businesses. A new project underway is evaluating the effect of a variety of major configuration changes on end-to-end (browser to mainframe) performance. The team has delivered several successful projects that demonstrate the value of PE services. The PE group is developing seminars and brown bags to advertise these services to the entire IT organization. The first of these was delivered in September 2000. Again, education and awareness are key to their success.

To get started, the group is providing performance modeling and analysis of several mainframe applications, the backend of the online services. The group has demonstrated to upper management that they can accurately predict the performance of core applications and identify improvements that should be made. Actual changes to the system based on these recommendations proved that the modeling and analysis had accurately predicted where the system bottlenecks would be and where IT should spend dollars on new hardware and software.

Additionally, the team is analyzing end-to-end configuration alternatives that would improve the performance of their Web site. This study helps business owners and the network performance group to understand the cost-effectiveness of making dramatic changes to the network and application infrastructure. A model of the current architecture was developed and proven accurate for predicting the response times measured in production. Changes to the architecture were added to the model to determine the effect on response time and performance. In just a few weeks, the performance engineers determined that the response time improvement that would be achieved did not have the magnitude to warrant the expense of altering the infrastructure. The company saved millions of dollars and months of effort. The project has been extended to evaluate other changes to the architecture that would have more significant impact on their response times. Several improvements have been identified that will provide much better return on investment for the IT department.

At this time, these are point solutions and successful projects that are not part of an integrated performance certification plan. As the team achieves its goals of accurately predicting performance of applications, it is also developing a performance plan for their entire online financial services system. This will include developing a plan for:

- Collecting performance data in a cohesive consistent manner,
- Building a framework model of the entire online system,
- Integrating data collection and modeling with the performance testing and the capacity planning organizations.

This organization will provide the same types of services as the insurance company's PE group except load testing. The financial group restricts their services to reviewing test plans and results, not performing the tests themselves.

The performance engineers are slowly moving projects earlier in the lifecycle chain to analyze the design of new applications. Providing the services of the PE group early gives the Financial Services company more time to respond to the recommendations that are delivered. A primary hurdle to delivering solutions early in the lifecycle is that system development schedules are extremely aggressive so development staffs are very difficult to get access to. However, with the support of the senior executives, the PE group is moving forward with plans to provide the services listed above. Already, they are demonstrating value by identifying end-to-end performance improvements that are cost-effective and ensure customer's access to online applications 24x7.

5.3 Key Ingredients to the Performance Engineering Practice

For both companies, the implementation of the performance engineering practice takes a conscious and dedicated effort. Both PE groups face the same challenges: 1) key application experts have very little time to support PE efforts, and 2) IT organizations are diverse and focused on delivery instead of performance. Both have learned that there are several very important elements needed to institutionalize performance engineering successfully:

- Obtain executive support to get the practice started.
- Build partnerships with both the technical organizations and business owners.
- Evangelize, educate and build trust throughout the IT organization.
- Guarantee success by delivering value to the business and technology partners.
- Reserve time for support from the application experts.

The result must be *zero time* added to the ***total*** system lifecycle by doing it right the first time.

6 Conclusion

The goal of performance certification is to mitigate risk in zero time. Innovation is required to develop a new organization that is focused, funded, sponsored and technologically prepared for performance engineering. With the right organizational structure and management commitment, no extra time need be added to the system development schedule. In fact, with reuse and integration of tools, methodologies and processes, the schedule can possibly be shortened. At present, performance can be actively planned, monitored, and verified throughout the system lifecycle. This requires a conscious effort to make cultural changes, process changes, and adopt new technologies. The value and benefit far outweigh the effort required to institutionalize performance engineering.

The creation of a performance plan and certification program provides a clear method for instituting change and verifying system performance at each step of the lifecycle. All of the involved organizations benefit from this incremental proof that the system is meeting business goals. The business owners get value and can calculate the ROI from the application features and functions. The IT department benefits by building a system that is cost-effective to implement and maintain. In the IT organization, the architecture, development, QA and operations groups benefit from performance certification. Each organization knows that the product delivered to them meets performance requirements. Thus, performance engineering improves inter-organizational communication, breaks down barriers, and builds a more cooperative team. Performance certification means certainty: each group receives a pledge that the deliverables at every step of the lifecycle have met performance requirements. Each group becomes committed to the process that insures performance.

eBusiness customers interact with entire IT systems, where some parts were once only accessible by employees. The performance of the system is more critical than ever before because the value includes customer retention, productivity and revenue. Institutionalization of performance plans is the only means by which the business owners can guarantee the experience of their customers, the responsiveness of their Internet applications, and the continued success of the business. The performance certification at the heart of the plan forms a solid basis for risk mitigation. When adopted into the culture of the system development lifecycle, performance engineering can be done in *zero time*.

7 References

[BART2000] Andrew Bartels, "The Difference Between e-business and e-commerce", Computerworld, October 30, 2000.

[LEWI2000] Ted Lewis, "Innovation Equals Risk Reduction", IEEE Computer, October 2000.

[NEUS1998] Douglas M. Neuse, "Why Simulate?", Capacity Management Review, February 1998.

[SMIT1990] Connie U. Smith, "Performance Engineering of Software Systems", Addison Wesley, 1990.