

# Understanding the Network Impact on Application Load Testing

A Shunra Software White Paper



## Introduction

Rigorous performance testing and optimization is a critical factor in the successful delivery of any business application. Yet frequently the performance of deployed applications doesn't live up to business requirements or end-user expectations. One reason behind these unpleasant "surprises" is the fact that most performance staging labs only test the application with local users (in a local area network (LAN) environment), while the fully deployed application is used by a variety of end-users, some local and others accessing the application remotely over different network links. The different network conditions that exist between end-users and application servers have a tremendous effect on the overall performance that remote end-users experience. This deviation from performance test results obtained in the lab is further exacerbated for N-Tier applications where each tier may reside in a different geographical location with it's own unique set of network conditions.

## Ensuring Application Performance to Remote End-Users

A simple example of such a deviation can be seen from the results of the following test conducted at a major North American financial institution. This test included both local users, (i.e. users interacting with the database over a local, unrestricted link), as well as remote end-users located in San Antonio, connecting to the datacenter in Tuttle, Ohio over a wide area network link. The wide area network and end-user load were created using emulation technologies. The emulated wide area link was a T1 connection, which accommodated the latency (geographic distance) between Columbus, OH and San Antonio, TX as well as the expected packet loss. Furthermore, the link was 30% utilized before the test in order to replicate the exact conditions in the production environment where the same network link supported other applications. This left 70% of the network link available for our financial application in the test below.

Location	Max Throughput (kbps) & Bandwidth Utilization (%)		Scripts		
	In	Out	Name	Max Load	Transaction Response Time
Local users in Tuttle, OH with local datacenter (LAN)	3838.29 (0%)	658.84 (0%)	Financial Transaction A	25	50.2
Remote end-users in San Antonio, TX	1168.66 (100%)	656.86 (80%)	Financial Transaction A	25	80.48

Figure 1. Summarized results of testing a financial application over a local network and wide area network.

The results in Figure 1 show a substantial deviation from the performance of the local end-users on the local area network to that experienced by end-users at the remote location on the wide area network link. While increasing the load of local end-users did not seem to strain the application server or backbone network, the same gradual load increase from the remote locations negatively affected performance due to overwhelming strain on the network and the application servers as well.

Now let's drill down further to understand the end-user behavior at each location and the causes of performance problems over the wide area network.

The graph in figure 2 below shows the performance of end-user transactions over time (red line), over a local network link, as the number of end-user transaction scripts running increases. In this case, the transaction response time over the local network increases only slightly, reaching 60 seconds at 25 end-user transactions (i.e. 33% increase from the starting point of 7 users at 45 seconds). This means that while the users are running over a local link, the system can scale and support additional end-user load.

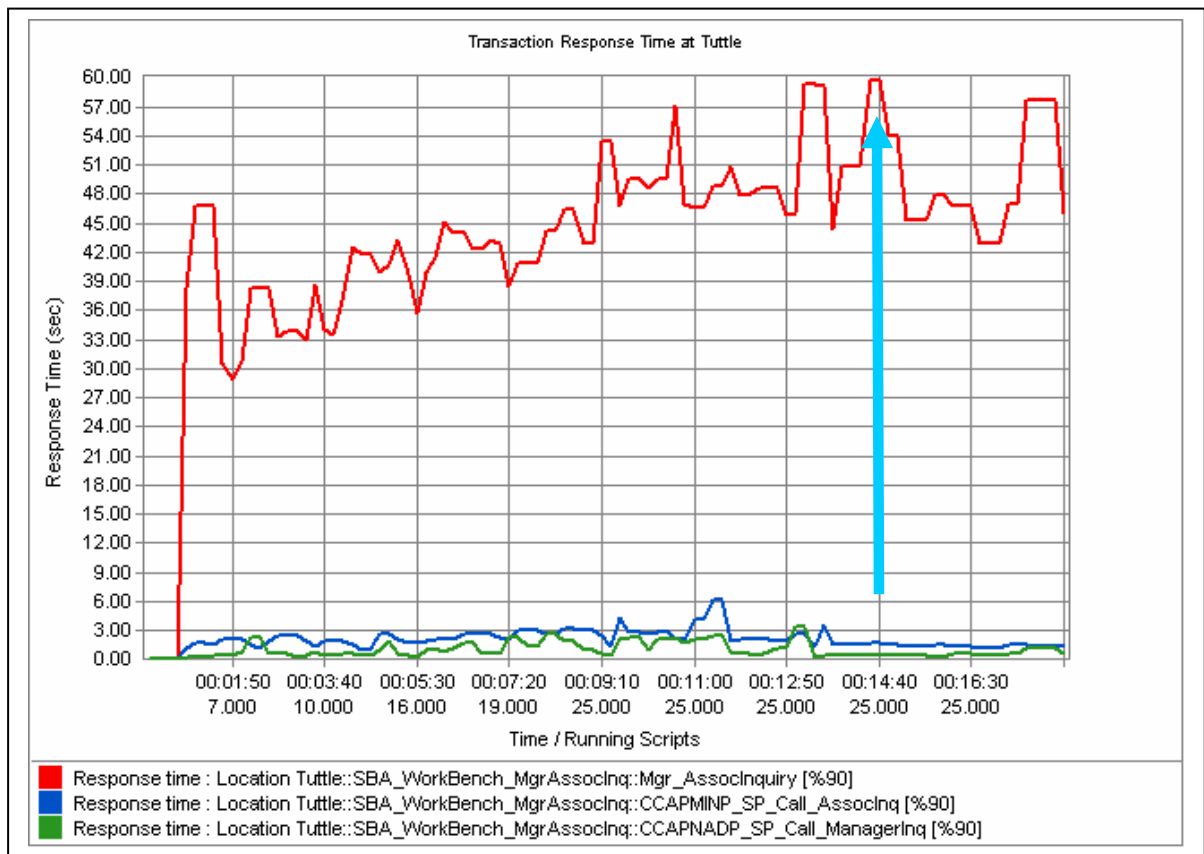


Figure 2. Results from testing end- user load over a local network.

However, when we ran the same test over an emulated wide area network the results were drastically different.

The graph in figure 3 below correlates end-user load with wide area network utilization. In this case, we can see that with as few as 16 end-user transactions (red line) running over an emulated wide area network, the available network bandwidth (light orange line) is already 90% utilized. Furthermore, at 25 end-user transactions the response time is up to 100 seconds (an increase of 66% over the response time over the local network link). Furthermore, nine minutes into the test, the application started generating errors (this is not shown on the chart). Obviously, this application is not ready for deployment to remote end-users over the wide area network.

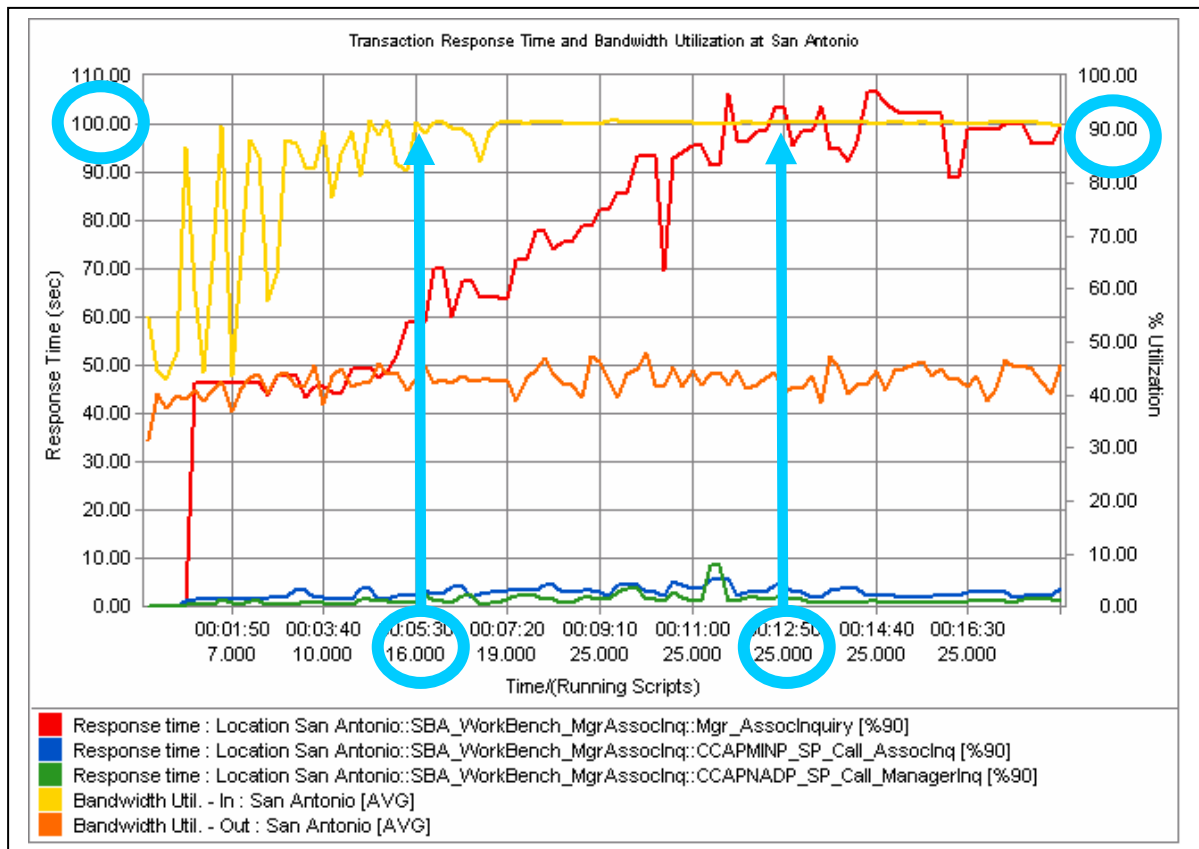


Figure 3. Results from testing end-user load over a wide area network.

## It's Not Just the Network

The tests above clearly showed that production wide area network conditions can significantly impact the performance of distributed applications, but it doesn't tell the whole story. In addition to the delay from the constrained network resources, the stress on the servers also increases since more resources are needed to support remote end-users. Sessions are open longer, OS resources are occupied for more time and more concurrent threads are needed. None of these factors are taken into account when testing and optimizing the performance only for local end-users on the local area network. Thus, remote business users are often unpleasantly "surprised" by poorly performing applications.

One of the primary causes of application performance shortfalls is insufficient hardware capacity. Any capacity planning that is done without considering the remote end-users is lacking vital information and may result in misleading or incorrect recommendations by ignoring the very real network and server conditions that exist on every corporate network.

## I Don't Control the Network – Why Should I Care?

Although the network appears to be someone else's responsibility, (Ah, so that's what those people in the network group do), no one can guarantee that the corporate network will provide acceptable conditions for optimal application performance for all end-user groups. If an application can't handle the conditions on the corporate network, there is very little anyone can do on the network side to fix it. It is really the responsibility of the entire development team, and QA in particular, to verify that an application can cope with the harsh conditions often found on the corporate network. Many solutions that improve the application's networkability rely upon 'tweaks' to the application itself rather than changes to the network. This is true regardless of whether it is in the way the client retrieves information (e.g. entry by entry Vs. retrieving bulks of information), the way the application is broken into different tiers (i.e. keeping the meta data closer to the user, vs. keeping all tiers in one geographical location) or many other best practices that dramatically improve the delivery and performance of an application to remote end-users.

## Can I Use My Load Tool's "Think Time" and Bandwidth Option to Create the Network Effects?

Adding "think time" to the end-user scripts, or using the bandwidth option of a load tool basically cause the virtual user's transactions to operate at a slower pace as they advance through the script. Simply adding a fixed transaction delay or artificially lowering bandwidth does not accurately represent real world conditions where remote end-user transactions cause contention for infrastructure resources, introduce bottlenecks along the connectivity path and experience delay due to the geographic distance and the load on the network. The "think time" option assumes a linear relationship between the effects of being remote and the response time of the application but the real world is much more chaotic and complex. The bandwidth option assumes that all end-users (or specific groups of end-users) will be

affected equally when, in fact, this is rarely the case under real network conditions. While both “think time” and bandwidth options can mimic some very limited network conditions, they lack the sophistication to adequately reflect real world conditions and are, therefore, of little use when testing remote user application delivery and performance.

As previously mentioned, yet another factor that needs to be taken into account is the differential way the application servers scale under remote end-user load. When the end-users are remote, servers need to maintain each session for a longer period of time, thus, utilizing operating system resources (sockets, threads, processes) for prolonged periods. QA tests that don’t take all of these factors into account will generate skewed results, similar to the deviation found between local users and remote users in the previous test.

### **Taking Remote Users into Account without Leaving the Local Lab**

Delivering an application without considering the target network is not sufficient for load testing. It puts remote end-users in the position of uncovering the performance implications of the network conditions. The solution was found by using an innovative approach to current load testing efforts that enabled users to incorporate the effects of the wide area network into the testing process, within the local lab. The effects of the remote end-users were achieved using the Shunra Virtual Enterprise. The Shunra Virtual Enterprise provides an exact replica of the production environment including remote end-users, remote offices, and the production network. It integrates with existing infrastructure and QA tools, such as Mercury’s LoadRunner® and Astra LoadTest®, and Segue’s SilkPerformer®, extending stress, functionality, scalability and benchmark testing solutions so that they include and account for the distributed nature of applications, throughout the QA process. With the Shunra Virtual Enterprise, users can quickly and easily find and resolve problems during testing, not after rollout, and avoid unnecessary and repetitive testing.

## About Shunra

Shunra provides solutions that empower organizations to address service level and performance concerns up front – before deployment. The Shunra Virtual Enterprise (Shunra VE) solution provides accurate, highly granular insight into how networked applications will function, perform and scale for remote end-users. It creates an exact replica of the production network environment, allowing IT professionals to safely develop, test and experiment with applications and infrastructure before rollout, and effectively plan for growth and change. With solutions tailored for networking and performance professionals, software developers, and quality assurance staff, Shunra VE facilitates collaboration across all IT disciplines – so IT organizations can quickly and more efficiently uncover and resolve problems before they impact the business. This results in more timely, higher quality and cost-efficient IT services, and the ability to “Deliver IT with Confidence”

### Solutions for Any Enterprise

More than 1500 customers, including hundreds of *Fortune* 1000 and Global *Forbes* 2000 organizations, from financial institutions to manufacturing companies, retail, energy, media companies, as well as independent hardware and software vendors and telecommunications service providers, have gained measurable returns from Shunra’s solutions. Among them are: 3M, Boeing, Cisco, Dow Chemical, EMC, FedEx, General Electric, General Motors, JPMorgan Chase, Kelly Services, Merrill Lynch, Motorola, Nestlé, Pitney Bowes, and Vodafone.

### Corporate Information

Shunra’s headquarters are located in New York City and Kfar Saba, Israel, with worldwide offices in Singapore, UK, The Netherlands and India. Shunra is also supported through a global network of channel partners.