

Document # 206148

Guaranteed WAN Application Performance

An Evaluation of the Ipanema WAN Optimization Solution



A white paper
commissioned by
Ipanema Technologies

T H E
TOLLY
G R O U P

White Paper

December 2006

Table of Contents

Before using this document you must agree to the terms of usage.
These terms are listed on the final page.

Executive Summary	4
Understanding Quality of Experience	6
Apdex Essentials	
MOS Explained	
Guaranteed Performance In Real World	9
Overview of Scenarios	
Scenario 1: Point to Point	
Scenario 2: Some to Any – Bandwidth Competition at the Destination	
Scenario 3: Some to Any – Cooperative Tele-Optimization	
Scenario 4: Any to Any – Competition at the Source and the Destination	
Methodology and Detailed Test Results	19
Apdex Formula	19
Displaying the Apdex Value	20
Detailed Test Results	21
Scenario 1	
Scenario 2	

Table of Contents

Before using this document you must agree to the terms of usage.
These terms are listed on the final page.

Scenario 3	
Scenario 4	
System Under Test	22
Ipanema System	
Test Bed Diagram	
Application Task Details	23
Web Browsing (HTTP)	
Exchange (MAPI)	
Windows File Access (CIFS)	
SSH Transaction	
VoIP stream (H.323 + G.711 payload)	
File Transfer (FTP)	
Streaming Service (MMS protocol)	
Appendix A: Ipanema Optimization Mechanisms	26

Executive Summary

Given the importance of application performance over often limited and expensive WAN circuits, it is no surprise that many different technologies are employed by various vendors to optimize performance over WAN links. It is possible to measure benefits in terms of bandwidth utilization and transaction times for given applications. But such an approach, while valid, might provide little insight into the overall quality of experience that end users encounter, especially when there is competition for bandwidth between multiple applications and multiple sites.

As businesses deepen their dependence on the Internet or private WANs to achieve their critical business objectives, guaranteeing performance of essential applications over the WAN is crucial for achieving business success.

In fact, with the advent of new types of networks and applications which consume more bandwidth and require different treatments, efficient usage of the limited resources is becoming a very challenging problem with which to cope. To make things worse, as the number of employees in branch offices increases steadily, naturally more traffic is requested across the WAN, which results in greater WAN bandwidth congestion than ever before. Historically, WAN capacity has not met the needs of most branch office users for various reasons. Many vendors have introduced various optimization solutions to improve application performance over the WAN. Ultimately, such WAN optimization solutions need to provide consistently acceptable end-user Quality of Experience (QoE) for business-critical applications under all circumstances. To do so, it is essential not just to optimize the WAN but to optimize intelligently by taking into account the needs of business-critical applications.

Ipanema Technologies' Ipanema System guaranteed application performance in congested WAN scenarios even when multiple applications and sites were competing for limited WAN resources.

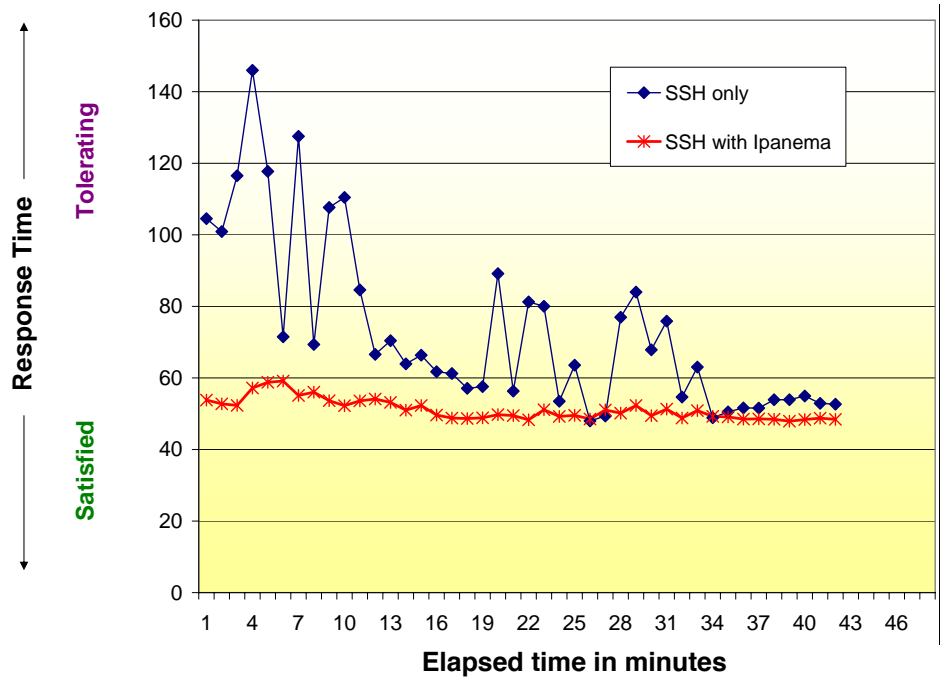
Ipanema Technologies, Inc. commissioned The Tolly Group to evaluate its Ipanema System Ver. 4.3 and illustrate its effectiveness in providing desirable and predictable application performance in congested networks. Tolly Group engineers subjected the solution to a battery of tests that involved running multiple applications in a simulated HQ/remote office three-node network. Application performance was benchmarked using a commercial-grade application simulation tool, Mercury Interactive's LoadRunner, in both

WHITE PAPER: Guaranteed WAN Application Performance

“before” and “after” scenarios. Results were summarized using an estimated MOS score for VoIP and the Apdex performance metric for all other applications. Apdex is a numerical measure of user satisfaction with the performance of enterprise applications. These metrics are better suited for capturing QoE than traditional response-time metrics (see Understanding Quality of Experience section for more details). Testing took place in October 2006.

Multiple test scenarios were run that measured the transaction times of common business-critical applications including VoIP, E-mail (MAPI), Web browsing (HTTP), SSH transactions and file access (CIFS) across a simulated 1 Mbps WAN link (with 50 ms of latency). Also, some non-critical traffic, such as FTP and Windows Media Streaming was used to introduce congestion into the WAN link for all the scenarios.

Response Time Comparison for SSH Traffic With and Without Ipanema Traffic Optimization



Note: The response times for SSH traffic optimized with the Ipanema System show slight variation and all fall within the “Satisfied” realm specified by the Apdex scoring system (discussed later in this report), leading to an excellent Apdex score. Without Ipanema, response times fluctuated markedly, leading to a poor Apdex score.

Source: The Tolly Group, October 2006

Figure 1

Ipanema
 Technologies,
 Inc.



Ipanema
 System Ver. 4.3

Network Optimization and
 Bandwidth Control

Product Specifications

Vendor-supplied information not necessarily verified by The Tolly Group

Ipanema Technologies
 Ipanema System

The Ipanema System delivers three main functions:

- Complete visibility of all application flows over the network
- Optimization of network resources through:
 - Dynamic bandwidth allocation
 - Smart packet forwarding
 - Advanced Compression
 - Adaptive TCP Acceleration
- Rightsizing: Optimal bandwidth sizing for access points according to application performance requirements

ip | e 5 appliance

- The new ip | e 5 appliance is a high-performing branch office device. It can handle asymmetric lines with up to 20 Mbps download, 1 Mbps upload and symmetric lines up to 8 Mbps download and upload.

For more information contact:

Ipanema Technologies, Inc.
 28 rue de la Pedoute
 92260 Fontenay-aux-Roses
 FRANCE
 Phone: +33 1 55 52 15 00
 Fax: +33 1 55 52 15 01
 E-mail: info@ipanematech.com
<http://www.ipanematech.com>

The device under test was optimized for each scenario and engineers utilized whatever optimization techniques available including Dynamic Bandwidth Allocation, Smart Packet Forwarding and Advanced Compression.

Tolly Group engineers measured the application response time or estimated MOS scores for each task for each application and calculated the Apdex score for each application. Mostly, there were less than a hundred samples collected for Apdex calculations. Then, Tolly Group engineers compared the Apdex scores of “before” and “after” scenarios (in other words, without and with the Ipanema product enabled) for the various business-critical applications for the different WAN scenarios.

Tests show that the Ipanema System, featuring the company’s ip | e 5 appliance at both ends of a WAN connection, guarantees the performance of business-critical applications by maintaining Excellent or Good Apdex scores (mostly Excellent) for all the business-critical applications across the range of scenarios tested. For the same tests, however, without the aid of the Ipanema system, testers were unable to guarantee the QoE, instead recording Apdex scores from Unacceptable to Excellent (mostly from Poor to Fair). (See Figure 1.)

Understanding Quality of Experience

Real-world business environments involve running multiple applications of varying importance across, usually, a variety of WAN links and locations. Different applications are more or less sensitive to different network scourges such as packet loss, delay and jitter. They have different end-user expectations; while receiving an E-mail with an attachment in 50 seconds could be acceptable, a transactional application such as SAP might require less than five seconds to complete, and a real-time application such as VoIP is even more demanding. In addition to the nature of the application — i.e. data-transfer, transactional or real time — its business criticality is an important factor to consider when it comes to the perceived QoE. It might be perfectly acceptable — even desirable — to provide average performance for the download of a Madonna clip from the Internet, while it would not be the case for the end-of-the-month SAP consolidation. The optimal QoE consists of offering the applica-

Apdex Alliance

The Apdex Alliance is a collection of companies collaborating to promote an application performance metric called Apdex. Apdex is a numerical measure of user satisfaction with the performance of enterprise applications, and reflects the effectiveness of IT investments in contributing to business objectives.

NetForecast

A collaborator for this test, NetForecast is an internationally recognized engineering consulting company that benchmarks, analyzes, and improves the performance of networked data, voice, and video applications. The company helps enterprises align application performance with business needs using a process based on the Apdex standard. NetForecast also advises technology vendors about customer requirements, technology issues, and the business value of application delivery products and services.

tion performance required to maximize the end-user productivity.

While metrics such as application acceleration — ratio of “before” and “after” response time — are easy to measure, they are not sufficient to reflect accurately the end-user QoE. A single response time value does not represent an end-user productivity and averaging several response times is not a good practice of measuring the QoE because the QoE has to be consistent. Also, an application accelerated by 30% does not tell us much about the end-user experience. Is 30% enough, or is 30% too much? Accelerating something that was already acceptable is not so meaningful from the QoE perspective.


Ultimately, it is beneficial to translate raw results into an index of overall user satisfaction, also referred to as “Quality of Experience”. The Apdex Alliance (<http://www.apdex.org>), a non-partisan industry organization, has created a tool with its Application Performance Index (Apdex). It converts measurements into one number on a uniform scale of 0-to-1 (0 = no users satisfied, 1 = all users satisfied). (See Figure 2.)

The Tolly Group used the Apdex scale exclusively to represent final results. NetForecast (a member of the Apdex Alliance, <http://www.netforecast.com>) assisted to ensure that Apdex was used correctly in this report. The following sections explain in greater detail the benchmarks used for this testing.

Apdex Scoring Breakdown

The Apdex Alliance has established the following correlation between numeric scores and descriptive ratings:

Excellent (Blue)	0.94 to 1.00
Good (Green)	0.85 to 0.94
Fair (Yellow)	0.70 and 0.85
Poor (Red)	0.50 to 0.70
Unacceptable (Gray)	0.50 or below



Source: Apdex Alliance, September 2005

Figure 2

Apdex Essentials

Apdex provides an effective way to distill complex and voluminous performance results into a numerical result that provides a straightforward basis for comparison. Customers evaluating WAN performance solutions will find comparison of Apdex values for competing solutions valuable.

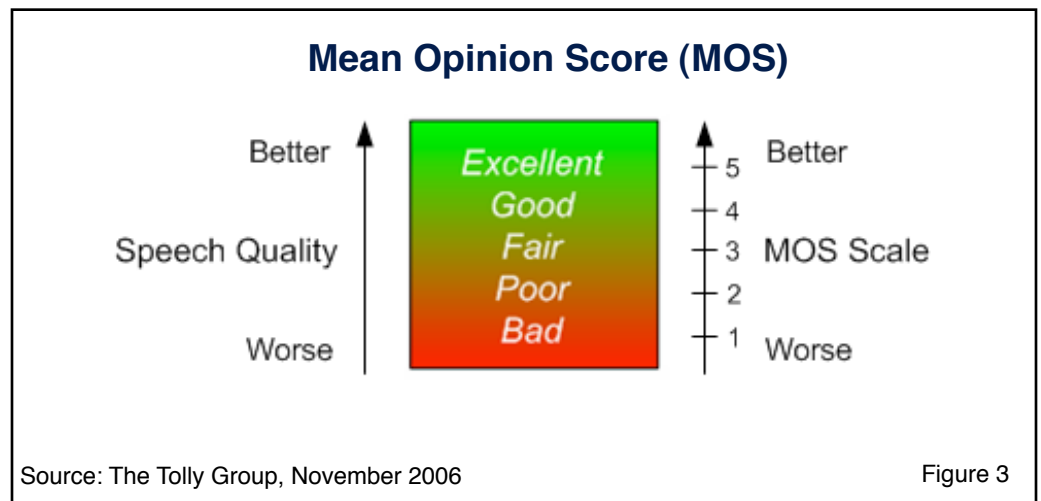
Apdex calculates a value based on the number of times a user's application performance was measured to be satisfactory, tolerable and/or unacceptable when compared with guidelines set by the business users organization. Apdex allows hundreds or more test runs to be distilled into a single, meaningful value.

Since Apdex values can be calculated separately for a variety of applications and then averaged, it is a useful measurement for aggregating data from complex tests involving multiple applications — essentially the conditions one encounters in the real world.

A perfect score for Apdex is 1.00 and it means that every application performed at an acceptable level in every measured sample.

MOS Explained

Mean Opinion Score (MOS) was designed by the ITU as a scale expressing PSTN voice quality based on subjective test involving real users. Users are asked to judge call quality among five different values from 1 (Poor) to 5 (Excellent). The concept has been extended to VoIP and objective testing using another ITU recommendation: E-Model. The E-Model takes as input a number of performance metrics of the VoIP



WHITE PAPER: Guaranteed WAN Application Performance

stream to compute an “R factor” which translates into an estimated MOS score. (See Figure 3.) In this test, we used the LoadRunner and H.323 Beacon, a tool by Ohio Supercomputer Center that emulates the VoIP stream and measured the estimated MOS score using the E-Model.

Guaranteed Performance In Real World

The main reasons for the deterioration of application performance are: the competition between applications, erratic surges in number of simultaneous users and, with the increasing deployment of MPLS networks, the meshed nature of the traffic. Testing the efficiency of WAN optimization technologies would fail to capture this reality if we use a “one link with one application” approach.

In the real world, virtually any Enterprise will need to support multiple applications simultaneously — often over a WAN configuration that is logically and physically a mesh — with applications traversing the various links in a peer fashion rather than an older style “hub and spoke” configuration where all branch offices communicated only with a central headquarters facility.

In four different congestion scenarios tested, the Ipanema System proved its robustness to guarantee mostly Excellent QoE for the business-critical applications and toll-quality VoIP delivery regardless of the types of congestion tested. In the same scenarios without the Ipanema System, however, the user experiences for the critical applications degraded and ranged from Unacceptable to Excellent (mostly from Poor to Fair) and the VoIP application could not even deliver business-quality voice.

Overview of Scenarios

In all of our tests, suites of applications were run both simultaneously and multiple times each — with varying amount of wait time — to create a more realistic, less “smooth” traffic scenario. Multiple simulated WAN scenarios were evaluated — each increasing in complexity, from “point to point” to “any to any” to reflect the meshed flow of traffic often encountered in the real world. As a result, there was competition between multiple applications and between multiple sites. Also, engineers emulated three different bandwidth competition scenarios: Competition at the source, competition at the destination and competition both at the source and at the destination.

Applications ranged from those critical in nature to those of a recrea-

WHITE PAPER: Guaranteed WAN Application Performance

tional nature and represented both interactive and file-transfer programs. See the “Application Task Detail” section for a detailed description of the applications emulated for the test. In any real Enterprise environment, the number of users of each application changes over time. In our simplified traffic model, the number of users of business-critical applications increased and then decreased successively throughout the period. For file-based transfers, (such as CIFS and Exchange) files were drawn from a pool of similarly sized files.

In every scenario, as in the real world, there were one or more “batch”, file-transfer applications. Such applications typically consist of a one-way feed of information and will typically use all of the available bandwidth until they have completed their task. VoIP and interactive transactions like Web browsing tend to exhibit degraded results in the presence of such high-bandwidth-demand applications.

Application response times were measured in “before” and “after” scenarios (i.e. without the Ipanema system active and with the Ipanema system enabled) for various WAN configurations. Tolly Group engineers emulated 1 Mbps and 50 ms one-way delay WAN using the Ipanema appliances to perform rate limiting and Netem for network impairment. (Netem is a Linux tool that emulates the properties of WANs. The current version emulates variable delay, loss, duplication and reordering of packets.) Then, the results were evaluated according to the Apdex process to relate the raw results directly to a QoE rating.

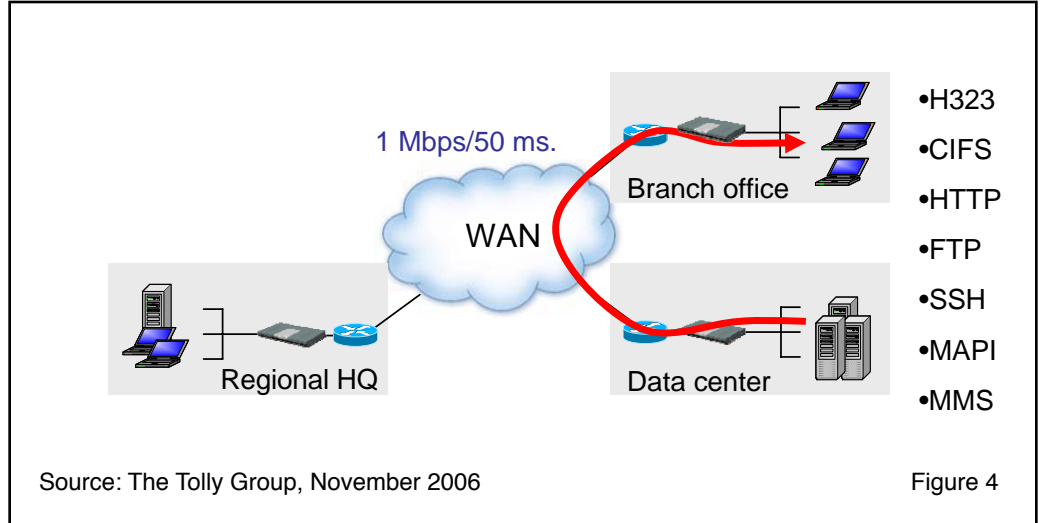
Determining acceptable performance is a challenge with applications. In the case of voice, though, there are several established methods for evaluating voice quality. One such method is MOS. The H.323 Beacon tool, integrated with Mercury Interactive’s LoadRunner, both used in this evaluation, can calculate MOS values based on end-to-end measurements made when running scripts that simulate voice traffic. Those measurements were used in this report.

In order to reflect the end-user QoE for other applications, engineers chose different target “T” values (Task response time) for different applications. The “T” values were defined based on NetForecast’s recommendation except for the SSH application since for this particular application we measured the time to complete the chain of tasks, not just a single task. The “T” values are based on what the users expect from business applications on a corporate networks.

Scenario 1: Point to Point

WHITE PAPER: Guaranteed WAN Application Performance

In this scenario, a set of branch office users were simulated accessing



the test suite of applications in a data center across an emulated 1 Mbps WAN connection. This scenario is called “Point-to-Point” since traffic flows in between a Data Center and a Branch Office (See Figure 4).

In the “Before” scenario, without the Ipanema system in place, the network was clearly overwhelmed by the demands of the user traffic. Only one of the applications even reached the Apdex level of 0.85 indicating “good” performance — the remaining applications experienced “Fair” or “Unacceptable” performance and the voice quality was unacceptably low by recording a MOS score of 1.87.

With the Ipanema system in place to prioritize and accelerate the traffic, every application improved dramatically. VoIP quality was measured at “toll quality” and application run times were all computed to “Excellent” Apdex scores. (Please see Figure 5 and refer to the previous “Apdex Essential” section for a detailed understanding of the test results.)

This dramatic improvement was due to Ipanema’s dynamic integration of optimization techniques with an intelligent real-time application monitoring capability. The optimization techniques that came into play in this scenario included Dynamic Bandwidth Allocation, Smart Packet Forwarding and Advanced Compression. These techniques are described in detail in later sections of this document.

Note that the detailed test methodology, test results and detail descriptions of each application flow can be found later in this paper in the section marked “Methodology and Detailed Results.”

QoE Index for Point-to-Point Scenario

Application	Criticality	QoE Index (MOS / Apdex)	
		Before	Ipanema
VoIP (MOS)	TOP	1.87	4.29
SSH	TOP	0.73 [64]*	1.00 [64]*
Web	TOP	0.42 [10]*	0.98 [10]*
CIFS	HIGH	1.00 [12]*	1.00 [12]*
Exchange	MED	0.73 [9]*	1.00 [9]*

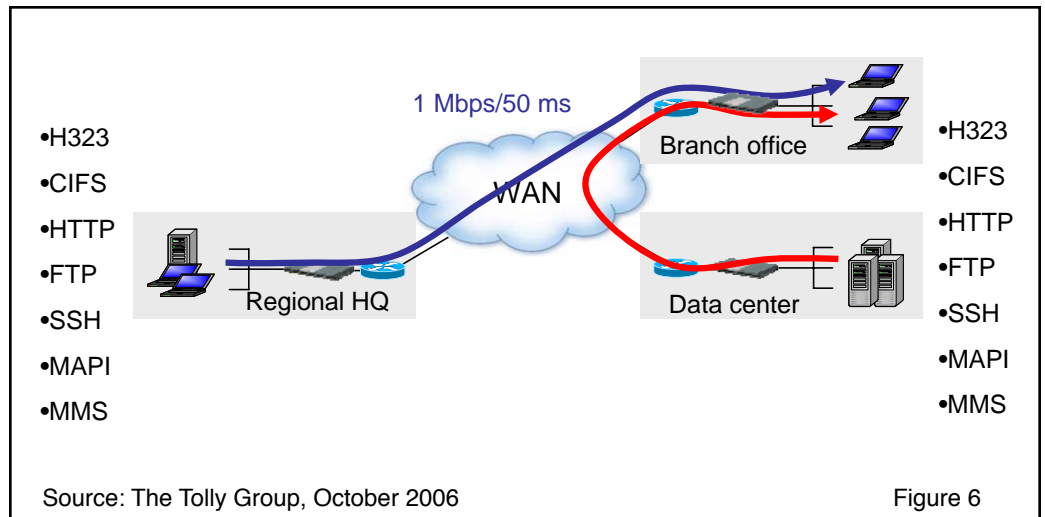


Note: Numbers inside square brackets [XX] represent the target response time on which the Apdex score is based. The asterisk (*) noted indicates that the Apdex score was computed with fewer than 100 samples.

Source: The Tolly Group, November 2006

Figure 5

Scenario 2: Some to Any — Bandwidth Competition at the Destination



In this scenario, a set of branch-office users were simulated accessing our test suite of applications in both data center and regional offices across an emulated 1 Mbps WAN connection so that major congestion occurred in the ingress queue of the branch office. This scenario is called “Some to Any” since traffic flows from some sites to any remote sites and results in bandwidth competition at the destination (See Figure 6.)

In the “Before” scenario, without the Ipanema system in place, the network was clearly overwhelmed by the demands of the user traffic. None of the applications even reached the Apdex level of 0.85 indicating “good” performance — most of them delivered “Fair” performance but a CIFS application registered “Poor” performance. The voice quality was relatively low by recording a MOS score of 3.12, which did not meet “business-quality” guidelines.

With the Ipanema system in place to prioritize and accelerate the traffic, every application improved dramatically in this scenario, too. The VoIP quality was again measured at “toll quality” and the application run times were all computed to “Excellent” Apdex scores. (See Figure 7.)

The same optimization techniques used in the Point-to-Point Scenario were utilized for this scenario. Note that the detail test methodology, test results and detail descriptions of each application flow can be found later in this paper in the Methodology and Detailed Results section.

QoE Index for Some-to-Any Scenario

		QoE Index (MOS / Apdex)	
Application	Criticality	Before	Ipanema
VoIP (MOS)	TOP	3.12	4.37
SSH	TOP	0.75 [64]*	1.00 [64]*
Web	TOP	0.74 [10]*	1.00 [10]*
CIFS	HIGH	0.60 [12]*	0.96 [12]*
Exchange	MED	0.83 [9]*	1.00 [9]*

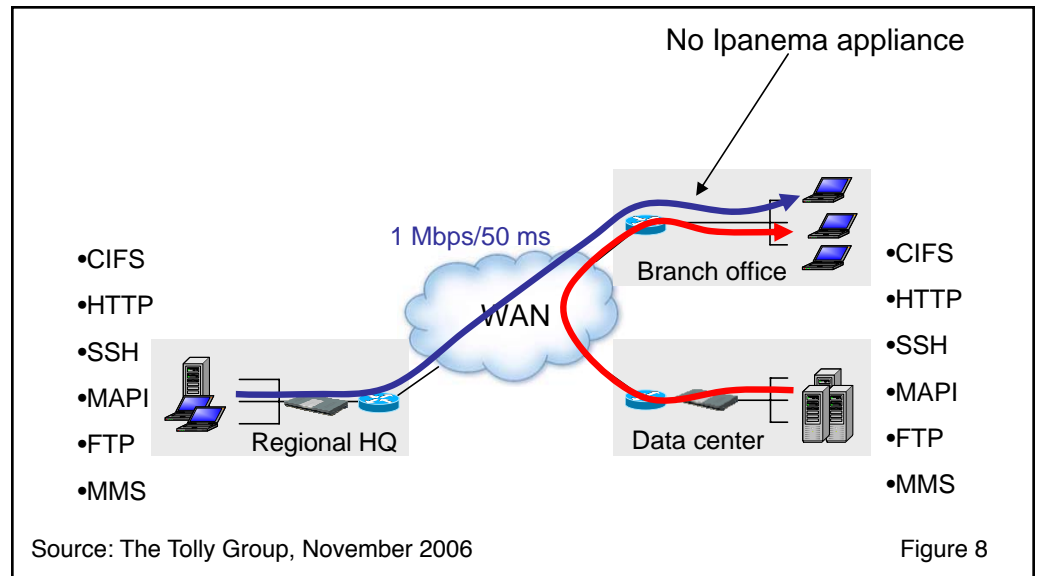


Note: Numbers inside square brackets [XX] represent the target response time on which the Apdex score is based. The asterisk (*) noted indicates that the Apdex score was computed with fewer than 100 samples.

Source: The Tolly Group, November 2006

Figure 7

Scenario 3: Some to Any — Cooperative Tele-Optimization



This scenario tested the Ipanema system’s patented Cooperative Tele-Optimization feature that allows users to control traffic headed for a branch office site without the benefit of a local Ipanema appliance at that branch location, thereby improving QoE at the site. This feature can be very attractive for large retail networks with many small branch offices receiving multiple flows from several data centers simultaneously.

The scenario had the same bandwidth competition between flows as Scenario 2 but with one major difference: there was no peer Ipanema appliance installed at the branch office. Cooperative Tele-Optimization can control the traffic going into the branch office in a collaborative effort between the data center and regional headquarters, thus being able to optimize competing flows from each source at the destination site. However, it cannot compress the traffic as an appliance would be required to decompress the flow at the other end. Finally, VoIP was not part of the application test suite as Ipanema recommend deploying devices on both end to fully control VoIP Quality of Experience. (See Figure 8.)

In the “Before” scenario, without the Ipanema system in place, results show that the network was clearly overwhelmed by the demands of the user traffic. Only one of the applications recorded “good” performance

WHITE PAPER: Guaranteed WAN Application Performance

– the rest delivered “Fair” or “Poor” performance and the voice stream was not generated for the estimated MOS measurement for this scenario.

With the Ipanema system in place to prioritize and accelerate the traffic, every application improved even though there was no peering device in the branch office. The performance of the business-critical applications still remained from “Excellent” to “Fair” (See Figure 9.)

QoE Index for Some-to-Any Cooperative Tele-optimization Scenario

		QoE Index (MOS / Apdex)	
Application	Criticality	Before	Ipanema
VoIP (MOS)	TOP	n/a	n/a
SSH	TOP	0.71 [64]*	0.85 [64]*
Web	TOP	0.82 [10]	0.95 [10]
CIFS	HIGH	0.65 [12]*	0.74 [12]*
Exchange	MED	0.89 [9]	0.98 [9]

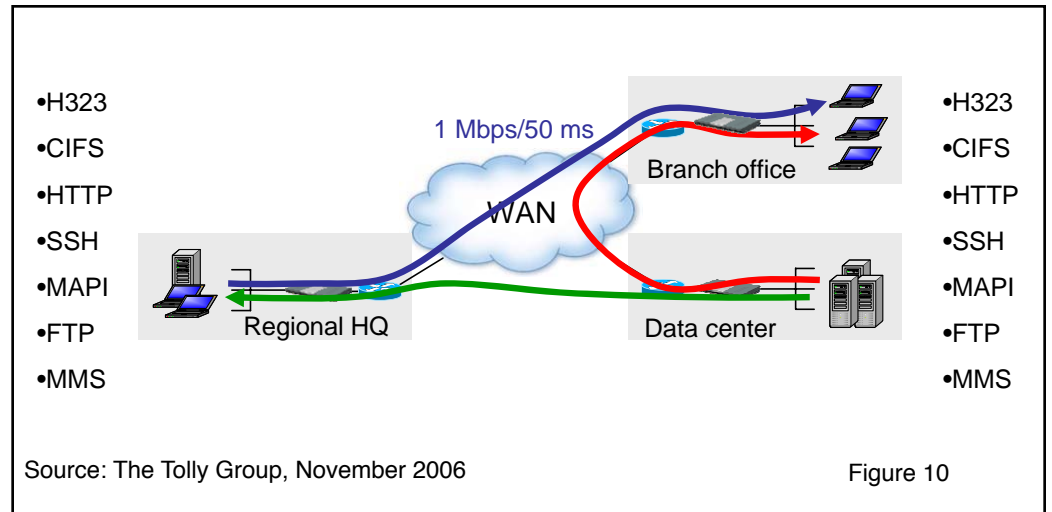


Note: Numbers inside square brackets [XX] represent the target response time on which the Apdex score is based. The asterisk (*) noted indicates that the Apdex score was computed with fewer than 100 samples.

Source: The Tolly Group, November 2006

Figure 9

Scenario 4: Any to Any — Competition at the Source and the Destination



The Any-to-Any scenario represented the most complicated scenario among the four scenarios tested. In this scenario, branch office users were simulated accessing test suite of applications in both Data Center and Regional HQ across an emulated WAN connection. At the same time, the Regional HQ was also emulated to access the applications in Data Center. This scenario created any-to-any full-mesh traffic flows, resulting in congestion at the source and the destination. (See Figure 10.)

In the “Before” scenario, without the Ipanema system in place, the network was clearly more overwhelmed by the demands of the user traffic than any other scenario. Overall “Before” performance in this scenario was poorer than the results from any other scenario. None of the applications even reached the Apdex level of 0.85 indicating “Good” performance — all of them measured “Fair” or “Poor” performance and the voice quality did not meet our “business-quality” guideline.

With the Ipanema system in place to prioritize and accelerate the traffic, every application improved dramatically. The VoIP quality was rated at “toll quality” and the application run times were such that most computed to “Excellent” Apdex scores. Only the CIFS application rated as a “Good” Apdex score (Please see Figure 11.)

This result proves that Ipanema systems would work well or even better in more complex full-mesh WAN scenarios. In other words, the QoE gap

WHITE PAPER: Guaranteed WAN Application Performance

in between “before” and “after” scenario would become large when the WAN network is complicated, such that different types of bandwidth competition among sites and applications occur at the source and the destination, which more closely reflects current Enterprise WAN environments.

Note that the detail test methodology, test results and detail descriptions of each application flow can be found later in this paper in the section marked “Methodology and Detailed Results.”

QoE Index for Any-to-Any Scenario

Application	Criticality	QoE Index (MOS / Apdex)	
		Before	Ipanema
VoIP (MOS)	TOP	3.17	4.28
SSH	TOP	0.65 [64]*	0.98 [64]*
Web	TOP	0.80 [10]*	1.00 [10]*
CIFS	HIGH	0.66 [12]*	0.89 [12]*
Exchange	MED	0.82 [9]*	0.97 [9]*



Note: Numbers inside square brackets [XX] represent the target response time on which the Apdex score is based. The asterisk (*) noted indicates that the Apdex score was computed with fewer than 100 samples.

Source: The Tolly Group, November 2006

Figure 11

Methodology and Detailed Test Results

Apdex Formula

Interested readers should consult the official Apdex specification which can be found at http://www.apdex.org/docs/Apdex_Technical_Specification.pdf

The Tolly Group followed this specification in calculating all Apdex values in this report with the assistance from the NetForecast. An Apdex value is always calculated, as per the specification, with respect to a target time established for a given application.

$$\text{Apdex}_T = \frac{\text{Satisfied count} + \frac{\text{Tolerating count}}{2}}{\text{Total samples}}$$

Apdex (Application Performance Index) uses the notion of QoE thresholds to calculate a single value representative of end-user QoE. This index is based on three zones of responsiveness:


- 👉 **Satisfied:** The user is fully productive. This represents the time value (T seconds) below which users are not impeded by application response time.
- 👉 **Tolerating:** The user notices performance lagging within responses greater than T, but continues the process.
- 👉 **Frustrated:** Performance with a response time greater than F seconds is unacceptable and users may abandon the process.

The Apdex formula is then the number of satisfied samples plus half of the tolerating samples, plus none of the frustrated samples, divided by all of the samples. For instance, if there are 100 samples with a target time of 3 seconds, where 60 samples are below 3 seconds, 30 are between 3 and 12 seconds, and the remaining 10 are above 12 seconds, the Apdex score is:

$$\frac{60 + \frac{30}{2}}{100} = 0.75$$

Displaying the Apdex Value

The Apdex values are decimal values between 0 and 1. The values always start with a 0, followed by a decimal point, followed by the fractional value for the calculation to two decimals. All Apdex values are calculated with a particular target threshold, T. The value of T must be clearly displayed in association with the Apdex score. When an Apdex value is the output of a sample of less than 100 instances, an asterisk (*) must be appended to that value. Thus an example of an Apdex value could be: 0.92[5.5]*



Source: Apdex Alliance, September 2005 Figure 12

Finally, the Apdex Alliance has established the following correlation between numeric scores and descriptive ratings:

Excellent	0.94 or above
Good	0.94 to 0.85
Fair	0.85 to 0.70
Poor	0.70 to 0.50
Unsatisfactory	0.50 or below

Detailed Test Results

Scenario 1

Application	Criticality	Satisfied MOS/ Application Response Time	Before			Ipanema		
			# Samples	Average MOS/ Application Response Time	MOS / Apdex	# Samples	Average MOS/ Application Response Time	MOS / Apdex
VoIP (MOS)	TOP	4.00	16	1.87	1.87	16	4.29	4.29
SSH	TOP	64.00	22	105.68	0.73 [64]*	24	51.97	1.00 [64]*
Web	TOP	10.00	17	39.05	0.42 [10]*	17	6.37	0.98 [10]*
CIFS	HIGH	12.00	10	3.42	1.00 [12]*	10	3.87	1.00 [12]*
Exchange	MED	9.00	20	15.49	0.73 [9]*	20	4.88	1.00 [9]*

Source: The Tolly Group, November 2006

Figure 13

Scenario 2

Application	Criticality	Satisfied MOS/ Application Response Time	Before			Ipanema		
			# Samples	Average MOS/ Application Response Time	MOS / Apdex	# Samples	Average MOS/ Application Response Time	MOS / Apdex
VoIP (MOS)	TOP	4.00	48	3.12	3.12	47	4.37	4.37
SSH	TOP	64.00	42	89.99	0.75 [64]*	42	50.65	1.00 [64]*
Web	TOP	10.00	30	15.44	0.74 [10]*	30	4.45	1.00 [10]*
CIFS	HIGH	12.00	21	27.74	0.60 [12]*	21	6.08	0.96 [12]*
Exchange	MED	9.00	59	8.48	0.83 [9]*	64	3.92	1.00 [9]*

Source: The Tolly Group, November 2006

Figure 14

Scenario 3

Application	Criticality	Satisfied MOS/ Application Response Time	Before			Ipanema		
			# Samples	Average MOS/ Application Response Time	MOS / Apdex	# Samples	Average MOS/ Application Response Time	MOS / Apdex
SSH	TOP	64.00	42	86.19	0.71 [64]*	42	59.42	0.85 [64]*
Web	TOP	10.00	103	12.21	0.82 [10]	103	6.26	0.95 [10]
CIFS	HIGH	12.00	40	23.26	0.65 [12]*	39	14.05	0.74 [12]*
Exchange	MED	9.00	180	6.68	0.89 [9]	180	4.22	0.98 [9]

Source: The Tolly Group, November 2006

Figure 15

Scenario 4

Application	Criticality	Satisfied MOS/ Application Response Time	Before			Ipanema		
			# Samples	Average MOS/ Application Response Time	MOS / Apdex	# Samples	Average MOS/ Application Response Time	MOS / Apdex
VoIP (MOS)	TOP	4.00	66	3.17	3.17	66	4.28	4.28
SSH	TOP	64.00	47	98.86	0.65 [64]*	48	53.95	0.98 [64]*
Web	TOP	10.00	52	11.81	0.80 [10]*	52	4.23	1.00 [10]*
CIFS	HIGH	12.00	26	24.75	0.66 [12]*	26	8.76	0.89 [12]*
Exchange	MED	9.00	58	10.69	0.82 [9]*	64	4.46	0.97 [9]*

Source: The Tolly Group, November 2006

Figure 16

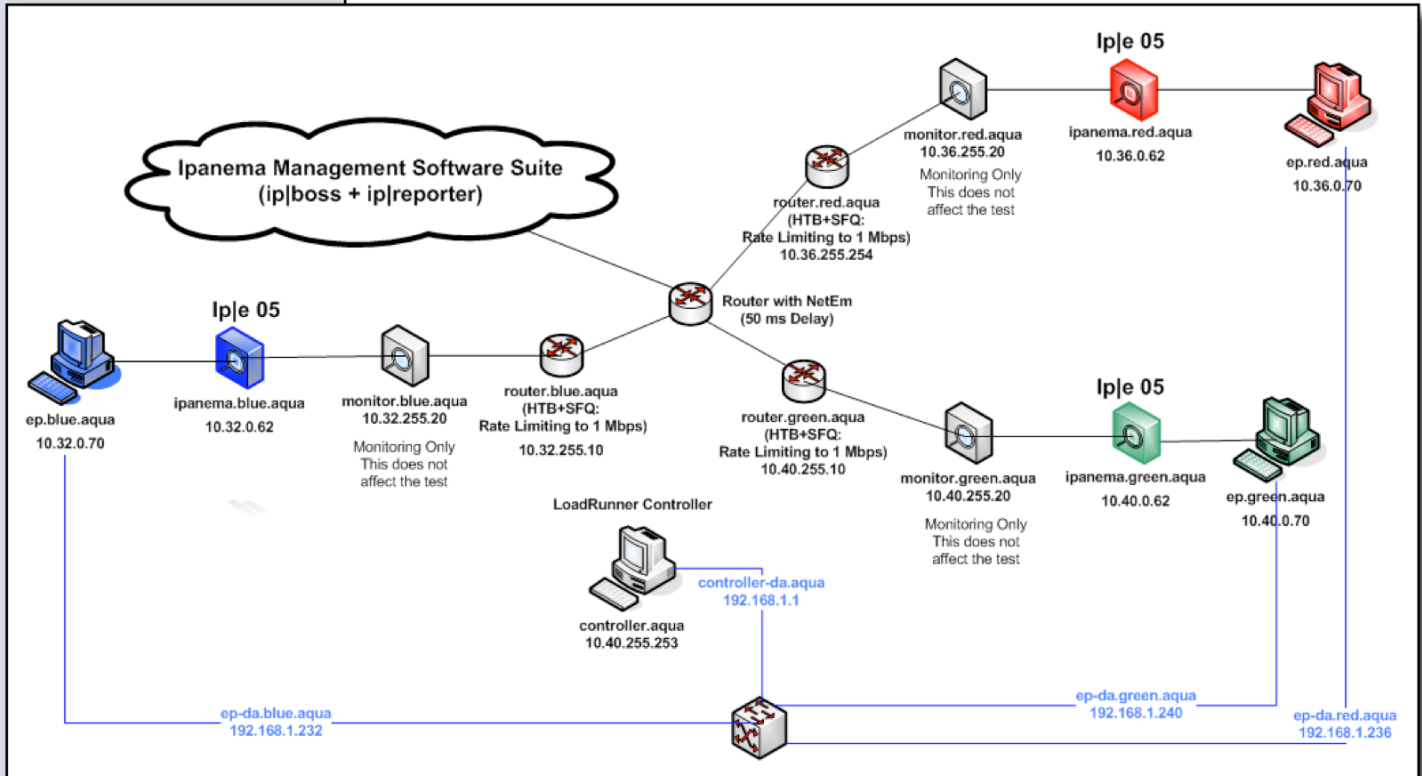
System Under Test

Ipanema System

Software component: Ipanema Management Software Suite V4.3 (ip | boss and ip | reporter)

Hardware component: ip | engine (Model: ip | e 5, V4.3 supporting 8 Mbps full-duplex WAN)

Test Bed Diagram



Source: The Tolly Group, November 2006

Figure 17

Application Task Details

Web Browsing (HTTP)

This task emulated browsing a PHP-based Web page. The Web page included one index PHP file, one Java Script file, one CSS file and 21 GIF files and the total Web page size is 115 KB. In this test, we considered business criticality of this application as TOP.

Exchange (MAPI)

This task emulated a Microsoft Exchange Server application implementing Messaging Application Programming Interface (MAPI). Each E-mail

WHITE PAPER: Guaranteed WAN Application Performance

user sent an message with an attachment. Attachment files changed randomly within a fixed pool of files for each task. The pool had 10 different files which range from 20 KB to 116 KB. In this test, we considered business criticality of this application as MEDIUM.

Windows File Access (CIFS)

This task emulated a remote file-system access protocol over the Internet, enabling groups of users to share documents across the Internet using the Common Internet File System (CIFS) protocol. Each emulated user reads a single file from the pool whose file sizes are from 150 KB to 300 KB. The accessed files changed randomly within a fixed pool of files for each task and the pool had 10 different files. In this test, we considered business criticality of this application as HIGH.

SSH Transaction

This task emulated a business applications based on ASCII terminal operations over SSH. The task was inspired from a real-world application involved in rental activities where the task starts with an ASCII screen that loads a number of forms which each end user is required to fill in. Then, the end user thinks for 10 seconds and starts filling the forms one by one (seven forms in total). The user also thinks one second between each form. The final step simulates pressing “confirmation” and receiving a results ASCII screen. The task was integrated into LoadRunner using a mix of batch and Cygwin (<http://www.cygwin.com>) scripts. This was the only application Tolly Group engineers did not measure the task time as per Apdex specifications. Instead, engineers measured the total time to complete the transactions mentioned above. In this test, we considered business criticality of this application as TOP.

VoIP stream (H.323 + G.711 payload)

This task emulated H.323 VoIP phone call using G.711u codec. The generated VoIP stream are real RTP (plus signaling) flows based on a 1.72 MB (8 KHz, 16-bit) wave file. The task was based on the H.323 Beacon Tool (<http://www.osc.edu/oarnet/itecoho.net/beacon/>). An estimated MOS calculation was performed using the H.323 Beacon Tool’s integrated functions. The task was also integrated into the LoadRunner test tool using a mix of batch and Cygwin scripts. In this test, we considered business criticality of this application as TOP.



WHITE PAPER: Guaranteed WAN Application Performance

File Transfer (FTP)

This task emulated a file transfer application using well-known FTP. Each task transfers a single file. Files changed randomly within a fixed pool of files for each task. The pool had 10 different files whose sizes are around 3.5 MB. In this test, we considered business criticality of this application as LOW and used for the Best-effort background traffic to introduce the congestion.

Streaming Service (MMS protocol)

This task emulated Microsoft Media Server (MMS) protocol and each emulated client downloads a media stream from the emulated server for 40 seconds at the rate of 16 Kbps. In this test, we considered business criticality of this application as LOW and used for the background streaming traffic to introduce the congestion.

Appendix A: Ipanema Optimization Mechanisms

The strong increase in QoE delivered by the Ipanema system is the result of three key characteristics of Ipanema optimization.

Ipanema Optimization is dynamic.

All the considered scenarios exhibit a high degree of variation in terms of number of users of the different applications. During a period, most of the WAN bandwidth can be utilized by, for example, the Web tasks with only a minimum number of users for other applications. Then, a few minutes later the dominant application can be VoIP, then MAPI and so on...

With traditional policy-based solutions, such dramatic changes would have required in most cases a manual modification of the parameters of the WAN optimization devices. This is especially true with all partition-based solutions. With Ipanema, those changes trigger an immediate, dynamic adaptation of the device's parameters. At each second, the system computes what are the optimal policies in each boxes according to the global performance objectives and according to a real-time analysis of the traffic mix.

Thanks to its dynamic properties the system is able to cope with the most difficult variations in terms of user demand leading to an always sustained appropriate QoE.

Ipanema Optimization is objective based.

The Apdex metric includes the notion of a minimum response time below which users are satisfied. This target response time can be achieved as long as the required resources are delivered by the network. In fact, a target response time can translate in most cases into a target minimum bandwidth per user as well as a maximum transit delay for packets for client-server applications, and also maximum jitter and loss for multimedia applications. Once the target response time has been delivered, i.e. the user has reached the maximum level of QoE, there is no need to throw in more of the precious network resources.

The Ipanema System does not require any understanding of traffic management policies to be configured. The only thing that needs to be defined



WHITE PAPER: Guaranteed WAN Application Performance

is the minimum set of resources the network should deliver so that end-users are happy. As a result, the system is able to take the target response times for an appropriate end-user QoE. Thanks to this objective-based approach, a much more efficient allocation of resource is implemented: what is optimized is, in effect, the QoE delivered to all users.

Ipanema optimization is global.

Modern WANs are based upon Multiprotocol Label Switching (MPLS) technology. As such, they benefit from a so-called any-to-any connectivity where each site can communicate with every site without having to go through a central location. This is truly revolutionary for WAN optimization mechanisms as they need to shift from managing competition between applications on a site, to managing competition between applications and between sites. Flows generating competition between sites to access the network resources are called meshed flows.

In real-life scenarios, the degree of flow meshing can vary. Many branches talking to a few data centers are creating meshed flows between the branches and the data centers. When branches talk together the meshing becomes more intense and the situation is described as "fully meshed." The Ipanema system takes into account meshed flows by delivering a global coordination between devices to obtain a full control of any application flow either in simple meshing situations up to fully meshed situations.

For more information see <http://www.ipanematech.com/>.



WHITE PAPER: Guaranteed WAN Application Performance

Terms of Usage

USE THIS DOCUMENT ONLY IF YOU AGREE TO THE TERMS LISTED HEREIN.

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase must be based on your own assessment of suitability.

This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions and certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks. Commercially reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur. In no event shall The Tolly Group be liable for damages of any kind including direct, indirect, special, incidental and consequential damages which may result from the use of information contained in this document

The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in Equivalent or better form to commercial customers.

When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from The Tolly Group's Web site.

All trademarks are the property of their respective owners.

The Tolly Group is a leading global provider of third-party validation services for vendors of IT products, components and services.

The company is based in Boca Raton, FL and can be reached by phone at (561) 391-5610, or via the Internet at <http://www.tolly.com>, sales@tolly.com



Entire Contents Copyright 2006 by
The Tolly Group, Inc.

ALL RIGHTS RESERVED