Tarantella® Enterprise 3™ Capacity and Scaling

IBM Intel Server Hardware
Red Hat 6.2 Operating System

A Tarantella White Paper **April, 2001**

Abstract

This white paper contains information useful for capacity planning for Tarantella Enterprise 3 servers. Tests were performed by Tarantella engineers under controlled conditions at IBM's RTP, NC Performance Lab. This paper describes the testing methodologies, results, analyses, and sizing guidelines for Tarantella Enterprise 3. The tests were done with IBM Intel Hardware and Red Hat 6.2 Operating System using Enterprise 3 version 3.01.

This white paper is for guidance only and may be subject to change.

Contents

Introduction

Tarantella® Enterprise 3™ server software leverages existing IT investment without the cost of re-engineering. It provides a non-intrusive solution that allows the enterprise to regain control and realize an unprecedented rate of savings. Enterprise 3 server software securely speeds access to Windows®, Web, Java®, mainframe, AS/400®, Linux® and UNIX® systems and applications from client devices anywhere in the world. This proven web-based solution's centralized management reduces complexity and effortlessly scales to accommodate rapid corporate change, technological advancement and expanding remote access needs.

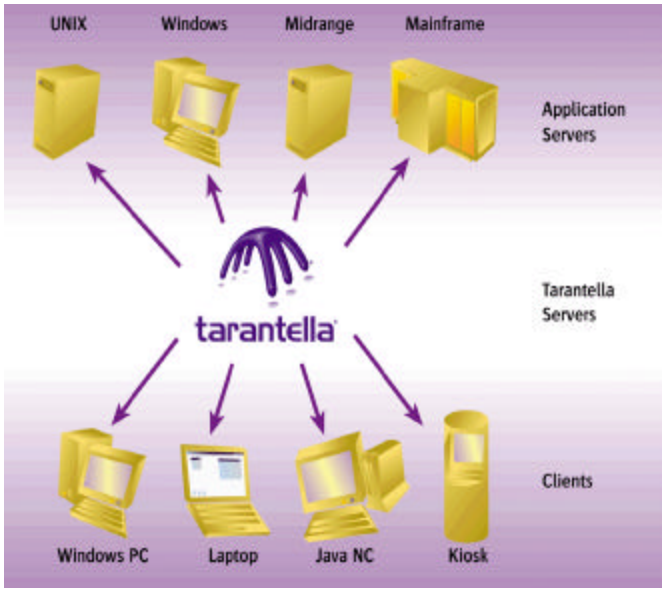

In a server-based computing environment, application execution and data processing occur on centralized servers. These servers can be any combination of Windows, Web, Linux, UNIX, or mainframe, servers. Tarantella software integrates and manages diverse server environments, providing a framework to deploy applications to users.

Information on Tarantella products, architecture, services, features, and licensing can be found at: http://www.tarantella.com.

This white paper contains testing methodologies, results, analyses, and sizing guidelines for Tarantella Enterprise 3.

To successfully deploy applications it is necessary to determine requirements for those servers. This can be determined by answering simple questions:

- What hardware is needed to support *n* users?
- How many users will a specific piece of hardware support?

The results and analyses contained within this document address these issues although they should not be interpreted in isolation. There are many factors to consider in a complex, server-based computing model. Sizing and capacity planning need to apply to all parts of the model, from the Tarantella server to the application servers being used.

Sizing and capacity planning information for Microsoft® Terminal Services can be found at: http://www.microsoft.com/windows2000/library/technologies/terminal/tscaling.asp

Findings Summary

Server capacity

The number of users that a server configuration can support varies depending on several criteria such as processor type, memory, hard disk, network configuration, and user type (typing speed, applications used, frequency, and so forth). The following table indicates the typical number of user sessions that a specific system can support.

In all cases, users were running Windows-based applications.

System ^	xSeries 230
Processor Type	Intel® Pentium® III
Processor Speed	1GHz
Processor(s)	2
Operating System	RedHat 6.2
Max user sessions tested	470
Recommended maximum number sessions	400

Table 1. Example Server capacity

System processor requirements

The recommendations outlined in Table 1 show the requirements for Tarantella web-enabling software and exclude the processor requirements of the operating system, which are minimal.

When running multiple applications, additional processing power may be needed. The amount depends on the applications being used. If two applications are being used and both are equally active, perhaps both generating reports, then twice the processing requirement is needed. However, if one application is being used while the other sits idle (as is more typical) then the only processing requirement is for the active application. In general, a requirement of approximately 5MHz of CPU performance for an Intel Pentium system is adequate.

Disk requirements

While disk capacity is not critical, we recommend selecting a large enough drive to handle the swap space needed (as determined by the amount of memory). We also recommend the use of fast hard drives to minimize any delays caused by swapping.

When running multiple applications the only impact on disk requirements is that each active, but idle, application is likely to end up in swap space. This means that the system swap space requirements may be higher.

System and user memory requirements

The recommendations in this document show the requirements for Tarantella web-enabling software and do not always detail the base requirements made by the underlying operating system. Enough memory must be allocated for the operating system; this can be 50-100MB.

As a general rule, memory requirements for Tarantella are approximately 128MB base, plus 4MB per active user session.

When running multiple application sessions more memory is required. However, some memory will be shared, and application usage patterns show that not all sessions are likely to be active at the same time. A simple rule to follow to adjust each user's requirement when running extra sessions is:

- For each extra UNIX X11 session, add 2MB per application.
- For each extra Windows session, add 3MB per application.

Tarantella arrays

A Tarantella array is a collection of connected Tarantella servers sharing a common set of configuration information, administered from a single point. An array acts as a single entity for administration and application deployment. The servers within a Tarantella array can automatically distribute users across the array to balance the load using a variety of load balancing algorithms. Naturally, adding more machines to a Tarantella array will increase the number of users that can be supported.

Adding more machines to a Tarantella array

Adding extra machines to a Tarantella array will always provide the ability to support more users. At a small array size the impact of multiple nodes causes minimal overhead, so total array capacity can be calculated by simply adding the capacity of individual arrays.

The overhead associated with array internal management will increase as the number of nodes and the number of user sessions increases. Generally, arrays of 10 to 20 nodes provide optimal resource usage.

Arrays versus multi-processor machines

There are two ways to increase the number of users supported in a Tarantella system: by adding multiple machines, and by adding multi-processor machines. Both have advantages, but the greatest factor when selecting the right combination is typically the cost of supporting more processors versus the cost of more, simpler machines:

- The use of Tarantella arrays can result in less expensive, less complex hardware that is easier to replace, should failure occur.
- The use of multi-processor machines can minimize the number of machines that must be purchased.

Tarantella arrays have a number of benefits, including automatic distribution of users across an array to balance the load and an added level of resilience should systems fail. Multiple machines can also offer better network capabilities, due to the additional network interfaces. However, there will be some additional network traffic within the array.

An array of multi-processor machines is recommended to gain the benefits of the array, while keeping the network traffic to a minimum.

Future expansion

While it is good to plan ahead and allow for the addition of extra processors and memory to Tarantella hardware, physical machines can always be added to the Tarantella system, at any time, as part of a Tarantella array. This allows for expansion at a rate controlled by the administrator and lessens the requirement to spend time and effort “up front” determining the size of machines. This avoids the necessity of purchasing new, large machines from day one. In addition, older machines don't have to be discarded as new ones are added, as Tarantella arrays may be built from a mix of system architectures.

Test Environment

Applications

Tarantella Enterprise 3 software can deploy many application types, from simple character-based to more complex and resource-consuming graphical applications. The tests covered by this paper concentrate on Windows applications using Microsoft RDP protocols. Windows applications generally provide the heaviest Tarantella load.

The type of application used has a large impact on resource usage, as does the way it is used. For example, a simple data entry application will take up less system and network resource than a large graphical application.

Users

The tests simulated the typical tasks of a "knowledge worker". Gartner defines these as workers who gather, add value to, and communicate information in a decision-support process. These resources are driven by projects and ad-hoc needs towards flexible tasks. Knowledge workers make their own decisions regarding what to work on and how to accomplish tasks.

Example job tasks include: marketing, project management, sales, desktop publishing, decision support, data mining, financial analysis, executive and supervisory management, design, and authoring. These workers create moderate to heavy usage patterns.

A user profile based on a reasonably intensive, interactive usage model was used. This uses an automated approach for driving Microsoft Excel via a scripted (parameter driven) macro.

The actual applications, features, and data sets used in these user scenarios cannot precisely mimic the experience of a real-life user on a moment-by-moment basis. This is because it is impossible to cover the wide range of tasks an individual is likely to perform within the confines of an automated test.

Hardware

Detailed descriptions of system configurations are provided later in this document. The basic configurations for the Tarantella server was dual Intel-based IBM® servers.

Software

The Tarantella Enterprise 3 server was running Red Hat 6.2.

Each application server was running Microsoft® Windows® 2000 Advanced Server with Terminal Services enabled and hosting Microsoft® Excel 97.

Each Client was running Microsoft® Windows®2000 with Internet Explorer.

Network

Since the aim of the tests was capacity planning, all machines were connected to the same network. No considerations were made for low-bandwidth connections.

Network specification:

Bandwidth: 100Mbit

Tarantella Server Kernel Tuning

The Tarantella server used in the tests was running Red Hat 6.2 with the following tuning.

The following lines were added to /etc/rc.d/rc.local:

```
echo "16384" > /proc/sys/fs/file-max
echo "49152" > /proc/sys/fs/inode-max
```

Running the Tests

The tests were intended to load the Tarantella server with users generating a level of stress that broadly equates to a typical "knowledge worker. Multiple browser sessions (Internet Explorer 5) were started interactively on each PC. Within each browser, a Tarantella user was logged in and the Excel macro was started and allowed to run continuously.

Application Parameters

An Excel macro was developed to test Tarantella software scalability with Windows applications. The macro loaded a worksheet, inputted data, produced a few graphs, and then reformated some of the data. It had delays throughout, to try to represent reasonably heavy usage.

The Excel macro generated a log file containing the time taken to execute each iteration. The macro had the following constants defined, allowing easy configuration of delay and of files accessed.

Excel macro constants (Note: all delays in milliseconds):

Delay between entry into cells	1000
Max delay between runs	2000
Min delay between runs	1000
Randomized delay	0 = off, 1 = on

Hardware

Tarantella Server Specifications:

Server type	^ xSeries 230	^ xSeries 230
Processor type	Intel Pentium III	Intel Pentium III
Processor speed	1 GHz	1 GHz
Processor(s)	1	2
Level 2 cache	256KB	256KB
Physical memory available	4096MB	4096MB

Client Specifications:

All clients used were identical. During tests 15-16 interactive sessions were run from each client.

Machine	IntelliStation M Pro™
Processor type	Intel Pentium III
Processor speed	500MHz
Memory	256MB
Network	100Mbit
Software	Windows® 2000 / Internet Explorer

Application Server Specifications:

Four Netfinity 7000 M10 systems were used as application servers.

Machine	Netfinity 7000 M10
Processor Type	Intel® Pentium II Xeon
Processor Speed	500 MHz
Processors	4
L2 Cache	2MB
Memory	4096MB
Network	100Mbit
Software	Windows® 2000 Advanced Server

Results

Single CPU Tarantella Server

^ xSeries 230

Results Summary:

Maximum Tarantella user sessions tested: 200

Memory used: 887MB

Memory per user: 3.1MB (256MB assumed used by system and Tarantella)

Processor usage per user: 5MHz

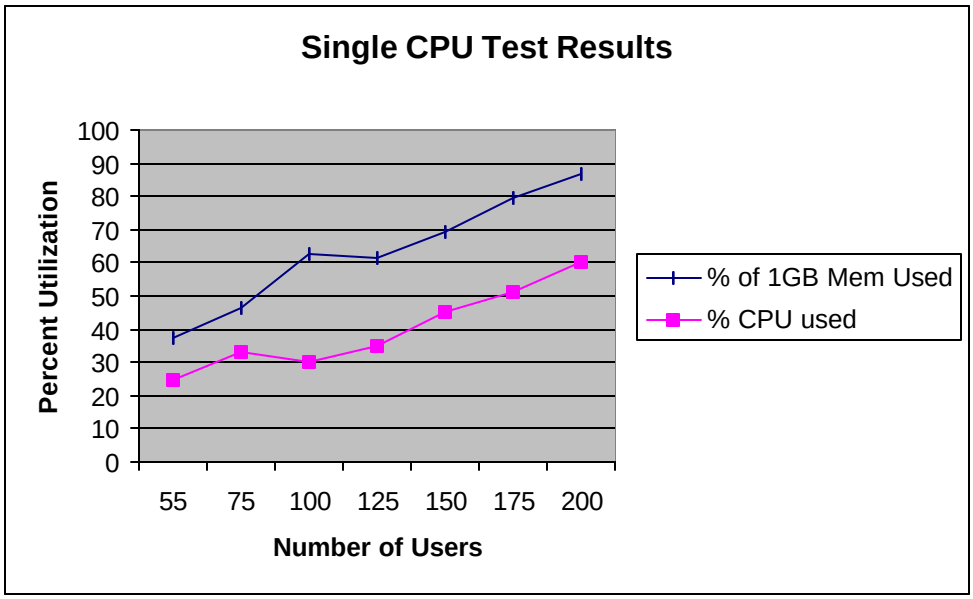

Dual CPU Tarantella Server

^ xSeries 230

Results Summary:

Maximum Tarantella user sessions tested: 470

Memory used: 1374MB

Memory per user: 2.4MB (256MB assumed used by system and Tarantella)

Processor usage per user: 5MHz

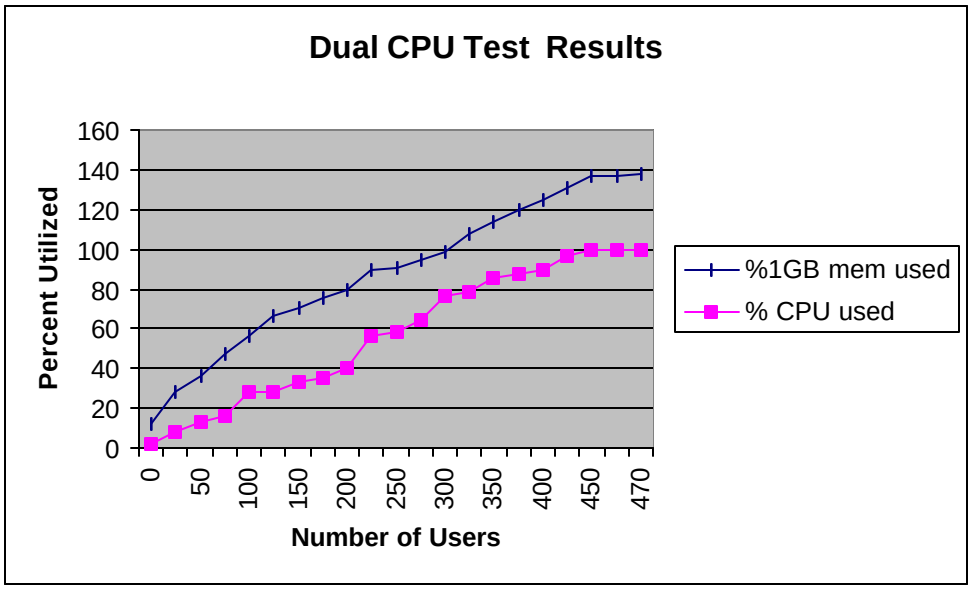

Conclusions

- Red Hat 6.2 is capable of supporting Tarantella Enterprise 3 on powerful Intel servers.
- Tarantella Enterprise 3 scales well on single-node systems and, with Red Hat 6.2, exploits Intel-based Symmetric Multiprocessor Systems well, with near-even distribution of load.
- ^ xSeries 230 system is capable of running up to around 200 users on a single CPU system and 400 users on a dual CPU system (depending upon the application mix involved).

Intel system memory and CPU requirements

The experience gained with sizing and scalability testing confirms the broad outline figures of 4MB memory per user, and 5MHz of processor for Pentium class systems.

However, determining the exact amount of memory the Tarantella server requires for a particular application is complex. It is possible to measure how much memory has been committed to Tarantella for each application the user starts—the memory the operating system has guaranteed Tarantella can access for each application instance. But the Tarantella server will not necessarily use all of that memory, and certainly is not using all of that memory at any one time. The subset of committed bytes that Tarantella has touched recently is referred to as the working set of that process. To complicate things further, the operating system can page this memory outside a process's working set to hard disk. The users and the applications are not affected by this but this should be taken into account when planning memory usage.

Performing Additional Capacity Planning Tests

To test or not to test?

The purpose of this document is to give a good starting point from which to make Tarantella server decisions. It may be more effective to go straight into pilot mode, rather than spend large amounts of resources analyzing users' work habits and capturing these actions into a simulated script.

Once a server configuration is chosen as a starting point (based on this paper's findings), users can gradually be added to determine the maximum number that a system configuration can support. If more server resources are required, it is always possible to add another server to a Tarantella array.

As an aid to understanding the various factors involved when deploying applications using Tarantella, the following should also be taken into consideration.

Determining application suitability

Some applications, such as those that make very extensive use of graphics or multimedia capabilities, may not be suited for running in a server-centric model.

If the application is suitable, it should preferably be run alone through the Tarantella server rather than starting an entire desktop session. This can save significant amounts of resources on the application server, and reduce resource requirements on the Tarantella server, thus allowing more users to log in simultaneously.

Characterization of users

Usage patterns need to be considered as they can have a significant performance impact on the application server, as well as the Tarantella server. For example, if all users log in at the same time of day, this will have an impact on the overall system responsiveness.

Network Utilization

Understanding the network environment is important with Tarantella. A Tarantella server can simply be added to an existing network to web-enable applications, However, an understanding of the network topology will always yield improvements in performance and scalability. Both network latency (the time it takes a packet to reach the other end of the network) and bandwidth (the amount of data that can travel over the network within a given period of time) are important factors. Because everything users see on their screens is server-based, both latency and bandwidth affect how the Tarantella server functions.

Tarantella web-enabling software makes use of the Adaptive Internet Protocol (AIP) to greatly reduce any problems that may be introduced by the network. AIP uses a number of methods to compress data and to remove redundant requests, which significantly help with low-bandwidth network connections. This makes the Tarantella server ideal for access from remote sites or for deployment of applications over the Internet, perhaps by an Application Service Provider (ASP). Connecting over a low-bandwidth connection has no significant impact on Tarantella server capacity or scaling.

High latency or load problems in other areas of the network, however, can have a negative effect on the responsiveness of any network. There is no way of solving these types of problems, and although they do not occur too often, they are unfortunately still encountered on the Internet, especially from more remote locations.

Although the Internet is an ideal and cost-effective network for business-to-business or remote site application deployment, we recommend that a dedicated network still be considered where the Internet causes too many latency or node stability problems.

When positioning Tarantella within a server environment some placement factors should be considered. Since AIP is designed to work over unknown and variable networks, the link from the Tarantella server to the clients should be the longest, while all connections between the Tarantella server and the application servers are minimized. So placing the Tarantella servers in the same room as the application servers they are connecting to is recommended.

Also, Tarantella web-enabling software makes use of the standard, native protocols that the application servers use, such as X11 and RDP. These protocols may not be designed to work over low-bandwidth networks or may not automatically adapt to the various network conditions that may be encountered. Therefore, a faster network between the Tarantella server and the application servers can yield better capacity and scalability. For example, increasing the network from 10Mbps to 100Mbps between these servers can result in improved performance and scalability. Note that it is not necessary to also increase the network connections between the Tarantella server and the clients, since AIP is quite capable of working over low-bandwidth connections. This can be a very cost-effective way to upgrade part of the network rather than upgrading both server and client connections all at once.

Further Information

For more information about Tarantella products, see the Tarantella web site at:

http://www.tarantella.com

For sizing and capacity planning information for Microsoft Terminal Services, see:

http://www.microsoft.com/windows2000/library/technologies/terminal/tscaling.asp

For information on IBM servers, visit:

http://www.ibm.com

Sales Offices

US & Worldwide HQ
Tarantella, Inc.
425 Encinal Street
Santa Cruz, CA 95061-1900
USA
Tel: +1 831 427 7222

Toll Free Sales and Information
Tel: +1 888 831 9700

www.tarantella.com
sales@tarantella.com

European HQ
Tarantella, Inc.
7 Brittania Court
The Green
West Drayton UB7 7PN
UK
Tel: +44 1923 816344

FreePhone Sales and Information
UK: 0800 0390134
France: 0800 913184
Germany: 0800 1802450
Italy: 0800 781920

www.tarantella.com
sales@tarantella.com

Pacific Rim HQ
Tarantella, Inc.
Nakamura Building, 2-4-3 Ohashi
Meguro-ku, Tokyo 153-0044
Japan
Tel: +81 3 5486 3924

Sales and Information
Tel: +81 3 5486 3924

www.tarantella.co.jp
japansales@tarantella.com