

*Scaling from 1 PC
to a super
computer using
Mascot*

*{MATRIX}
{SCIENCE}*

Mascot - options for high throughput

- **When is more processing power required?**

Options:

- **Multiple Mascot installations**
- **Using a large Unix server**
- **Clusters - hardware and software configurations**

{MATRIX}
{SCIENCE}

In this session I am going to talk about the three main options available if you need higher throughput for Mascot searches than is available on a standard PC.

So firstly I will describe how to determine if you need more processing power and if so how much. Then I will discuss the three options:

- Multiple installations, which we don't generally recommend
- using a large Unix server
- using a cluster

When is more processing power required

- For high throughput and quick response
- For large no-enzyme searches, or searches against a number of very large databases from a single instrument
- For multiple instruments

{MATRIX}
{SCIENCE}

So firstly, why might you need more processing power. You may think that this is all very obvious, but I can assure you that we get a fairly equal number of requests for quotes for overpowered and underpowered systems.

The main reasons are listed here:

Clearly if you have over 50 instruments running in a high throughput environment like ...



GeneProt, then a single PC running Windows is not going to be up to the job.

Worst case conditions

- **Large sequence database**
 - human genome
 - dbEST
- **Large MS dataset**
 - 10,000 - 100,000 spectra
- **Inaccurate / noisy data**
- **Large numbers of P-T modifications**
- **No enzyme specificity**

{MATRIX}
{SCIENCE}

So what are the worst case conditions -

The human genome has approximately 3 Gigabases, and the complete dbEST is now over twice that size. Searching these databases is clearly going to be extremely computationally intensive.

However, even a single instrument running a MUDPIT experiment can produce in excess of 10,000 spectra.

Inaccurate or noisy data can mean using large tolerances, and this will affect the search speed.

Large numbers of post translational or chemical modifications (selected as variable modification) will also increase the search speed. There is no time penalty for fixed modifications.

Lastly, a search with no enzyme specificity will dramatically increase search times.

How many CPUs do I need?

- **It is possible to make a reasonably accurate estimate if you have enough information about the types of search you want to perform . . .**
- **Search speed is almost exactly proportional to the number of CPUs and CPU speed**
- **Try a search on Matrix Science public web site at weekend**
- **Send us data and ask for timing - please ask first!**

{MATRIX}
{SCIENCE}

Estimating the number of CPU's required to achieve a given throughput is complex. One option is to try typical searches on our web site.

Using multiple Mascot installations

- **One Mascot server per instrument**

but...

- **More administration required - e.g. for database updates, backups of data, configuration changes etc.**
- **Cannot perform a very large search rapidly**
- **Harder to find old results files - need to remember which Mascot server was used**

{MATRIX}
{SCIENCE}

So the first, and most obvious solution is to have multiple Mascot installations - for example one per instrument, or one per user.

However, we would generally recommend against this because...

Windows / Linux 'SMP' systems

- **Windows 2000 Professional, XP and NT 4 Workstation support 2 processors**
- **Windows 2000 / NT4 Server for > 2 processors**
- **Linux smp**
- **4 processor Intel Box expensive - e.g \$30k for a 4 processor 1.6 GHz PowerEdge 6650**
- **2Gb memory limit. . .**

{MATRIX}
{SCIENCE}

For a simple Intel or Athlon based PC, the maximum number of processors that we would recommend is 4.

With Windows 2000 professional, XP and NT4, the limit is 2 processors.

The big Unix Box solution



COMPAQ hp & compaq

AlphaServer GS320 SYSTEM

- Up to 32 1 GHz Alpha processors
- Up to 256 GB of ECC memory



pSeries 690

8- to 32-way
Clock rates
(Min/Max)
1.1 GHz / 1.3 GHz
System memory
(Std/Max)
8GB / 256GB



SGI Origin 3000

Processors	2-512
System bandwidth	Up to 716 GB/s
Maximum memory	1TB



Sun Fire 15K Server

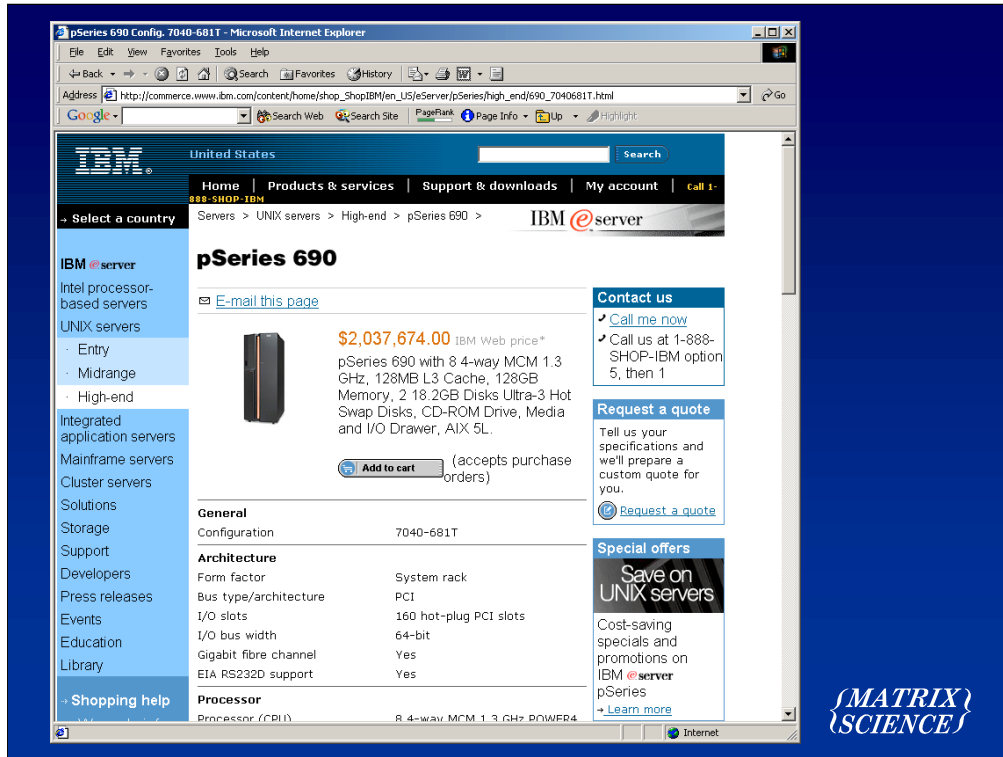
With 106 UltraSPARC® III Cu 900-MHz processors, more than half a terabyte of memory

{MATRIX}
{SCIENCE}

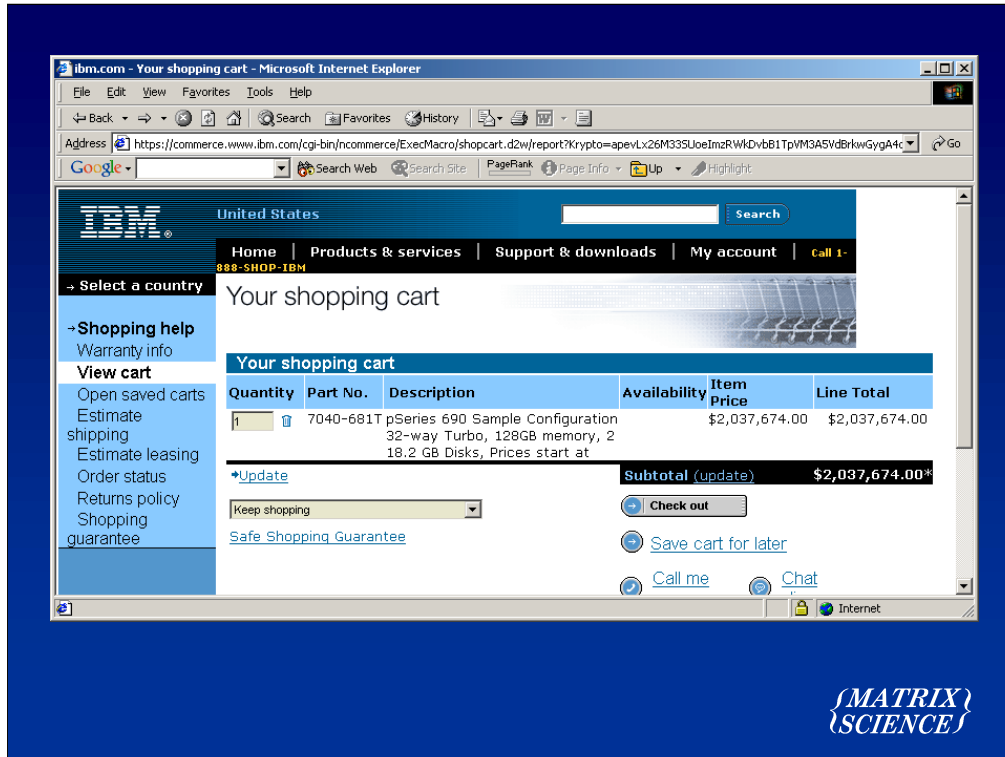
So then we have the big Unix Box solution. The four main manufacturers are shown here, and each system has its merits.

The advantage of these systems is that they really can scale.

...

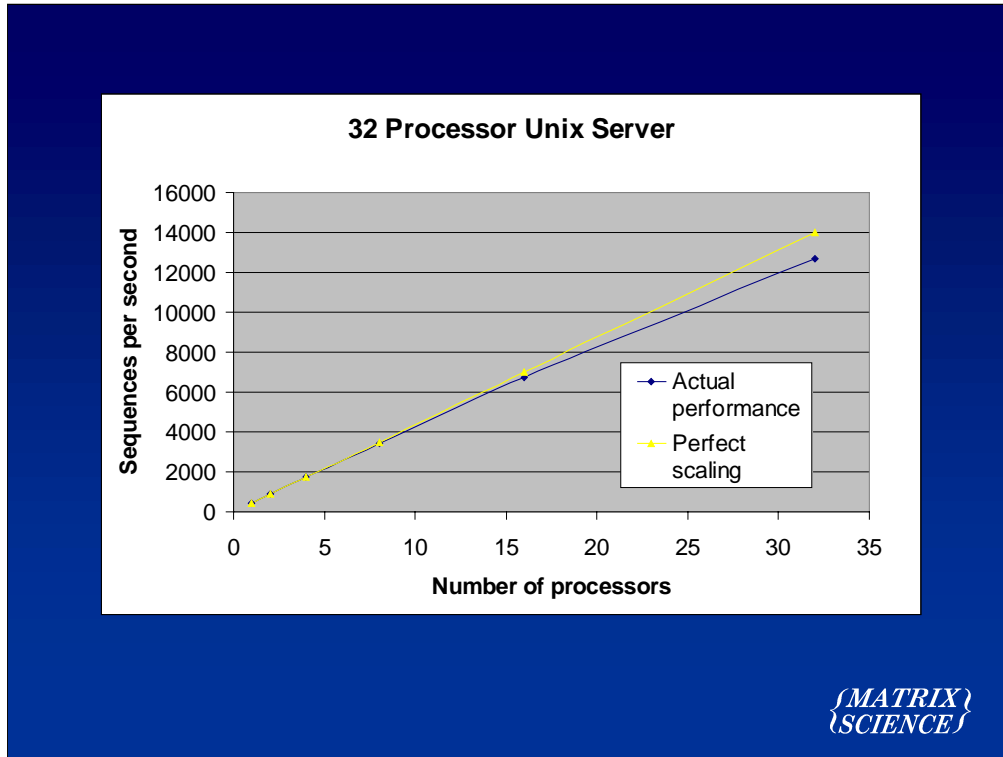


And what's more, they can now easily be purchased over the internet



{MATRIX}
{SCIENCE}

Anybody got a credit card I can borrow!



This graph shows how well one of these systems can scale. This graph is for a 32 processors Unix box, and the yellow line is showing perfect scaling. Once we get up to about 32 processors, it starts to tail off a little.

Advantages of a large Unix system

- **Often 'company' fasta files compiled on a large Unix box**
- **Robust / excellent support / well maintained**
- **Some systems can scale to more processors**
- **In house support from IT / Bioinformatics**
- **Large amount of memory**
- **Easier to maintain and support than a cluster**

{MATRIX}
{SCIENCE}

Disadvantages of a large Unix box

- Older systems can be slow - 2 year lifetime
- Price

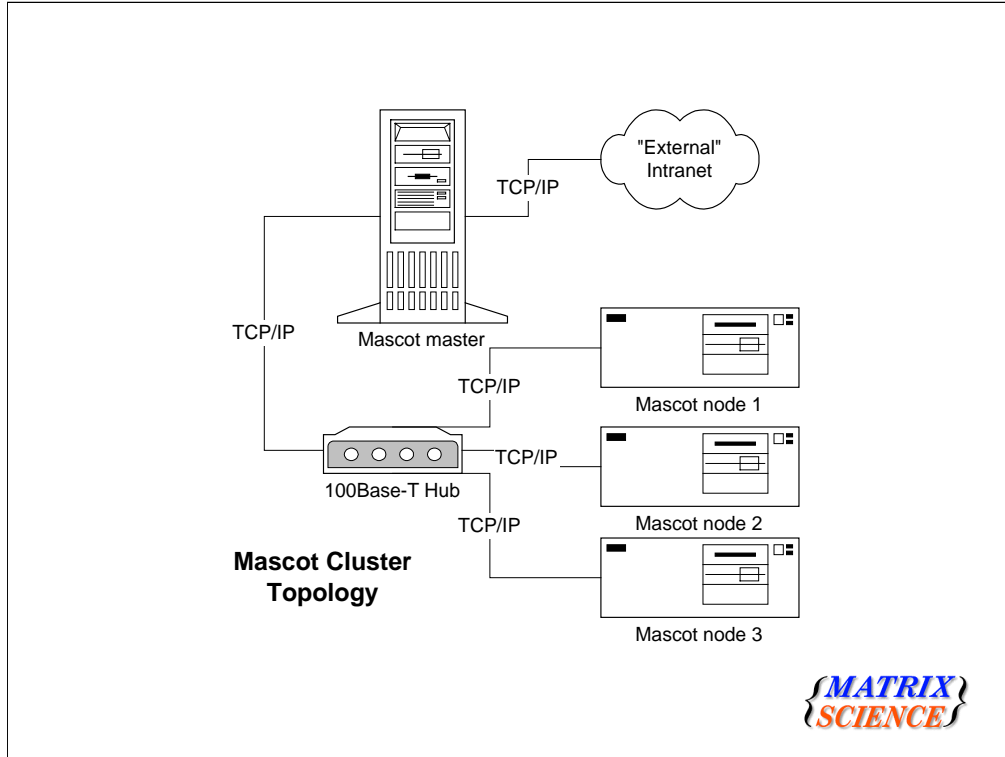
{MATRIX}
{SCIENCE}

A cheaper way?

- **Best value for money computers are dual processor systems**
- **A cluster of single or dual processor computers is a cost effective way of increasing throughput**

{MATRIX}
{SCIENCE}

A cluster of standard PC's is a very cost effective and scalable solution



Mascot supports cluster operation using a Beowulf-like topology. Mascot supports cluster mode as standard, whenever the licence is for 4 CPU's or more. You just have to hook up an appropriate number of PC's on a local LAN

Matrix Science Web server

- **19" rack-mountable chassis containing 10 fast Intel processors**
- **Redundant power supplies**
- **We run Windows on the nodes, and the 'Master computer' runs FreeBSD**
- **2Gb RAM per node**
- **Hot-swappable**
- **Similar modules available for sale as a turn-key package**



Alternatively, we can supply a complete turn-key solution based on rack mounted, high-reliability PC hardware. Our web server uses this hardware

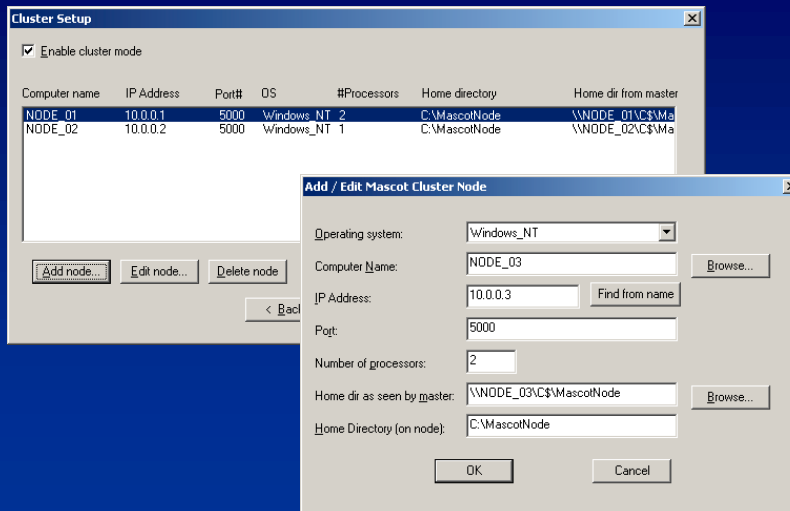
Software setup - Windows

- **Install Windows 2000 workstation on each node**
- **Install Windows 2000 server on master**
- **Run Mascot installation:**

{MATRIX}
{SCIENCE}

Setup under Windows is extremely simple

Cluster software setup - Windows



{MATRIX}
{SCIENCE}

If the licence is for 4 CPU's or more, you will be presented with additional dialogs for configuring the cluster nodes

Cluster software setup - Unix

- **Install Unix version of Mascot**
- **Edit nodelist.txt and a couple of scripts**
- **Applications and configuration files are all transferred to each of the nodes automatically**

{MATRIX}
{SCIENCE}

Under Unix, its just a case of editing two configuration files

Functionality of Mascot cluster

- **Software automatically installed onto each node - no need to 'pre-load' each system**
- **Sequence databases and all configuration files distributed / updated automatically**
- **Each search runs on every node so there is an (almost) linear increase in search speed**
- **'Heart beat' health check**
- **Automatic recovery if a node fails**
- **"Spare" nodes are supported**
- **Distributed memory**
- **System errors emailed to administrator**

{MATRIX}
{SCIENCE}

Cluster support includes many features to support large, high throughput systems

Status report

Node	IP Address	OS	Responding	Physical Memory	Swap file	Disk space
slave01	10.0.0.1	Windows NT	☺ OK	☺ 44% free	☺ 97% free	☺ 90% free
slave02	10.0.0.2	Windows NT	☺ OK	☺ 45% free	☺ 97% free	☺ 90% free
slave03	10.0.0.3	Windows NT	☺ OK	☺ 50% free	☺ 97% free	☺ 90% free
slave04	10.0.0.4	Windows NT	☺ OK	☺ 47% free	☺ 97% free	☺ 83% free
slave05	10.0.0.5	Windows NT	☺ OK	☺ 47% free	☺ 97% free	☺ 90% free

{MATRIX}
{SCIENCE}


"At a glance" status screen

What are the limits. . .

- **Less than 4 processors makes little sense**
- **Limit of 1024 CPUs per 'sub-cluster'**
- **With Mascot version 1.8, practical limit of about 100 dual processor 32 bit nodes**
- **Limit of 4096 nodes in total**

{MATRIX}
{SCIENCE}


Of course, there have to be some limits



GeneProt have the largest Mascot cluster. 1440 Alpha CPU's configured as a number of sub-clusters

MASCOT®

Take the guesswork out of protein identification...



MATRIX
SCIENCE
www.matrixscience.com