Sizing Guidelines for CleverPath Portal

The following sizing guidelines have been put together as the basis of a one hour techtalk for FSG presales to assist in sizing portal servers for new business opportunities. While there a number of types of portal solutions ranging from Business to Business (B2B), Business to Customer (B2C), and Business to Employee (B2E), the key thing will be to estimate the user loading for the portal solution. For illustrative purposes, and for what is the most common business situation we are running into, we will go thru the details of estimating loading for a B2E intranet deployment.

The latter parts of this document, starting with the section starting with "Using a Sizing Calculator", are almost self standing and can be used as a handy reference document for sizing. It ends with a spread sheet calculator.

We consider this a "living document" that will constantly be updated as more information becomes available. As an example substantial changes have been made to the calculator since the May techtalk.

Sizing Portal installations

Designing a Portal installation can be a difficult task, even though the network and system designs are usually quite simple. For many sites, the most difficult part of installing Portal is determining how many servers are needed and what size each one should be. To do this, we need to understand how expected use can be used to predict the resources needed, and how to convert the resource prediction into an acceptable system design.

Method

Our fundamental assumption is that Portal capacity is determined by the system's computing power. This assumption is key: we use processor ratings to predict performance for untested servers, for example. It also guides our design approach. This method will predict the same capacity for two servers, regardless of other factors such as their I/O capacity and storage speed, if they have the same type and number of processors. In practice, this may not always be true, but most differences between system designs will be small when compared to the uncertainties implicit in predicting usage, for example.

This assumption is based on practical experience. When the Portal server is benchmarked, there are few disk accesses, so disk performance is not a significant issue. Total amount of information transferred from the Portal server to the test clients is significantly less than the amount that could be accommodated by the system and its network. As a result, as well as our practical experience with the Portal product, we feel that the amount of computing power available to the site will determine how fast the Portal can operate.

This paper outlines a simple method for estimating the computing power needed for a specific installation. Given a consistent approach to estimating the expected load for the

Portal installation, and a rating for Portal performance, we can predict the computing power needed to satisfy the installation's user. The estimated power can then be used to select the specific hardware needed for the site's new Portal installation.

Limitations

There are a number of error and uncertainty sources in this process. While we can predict the number of standard workload operations that a given server can perform per second or hour, correlating this with the number of "real" pages that can be delivered would require detailed Web page design details to produce an accurate prediction. Predicting the capacity required to satisfy a large and distributed audience is an imprecise operation as well – Portal performance will have a direct effect on the users' attitude towards the system, which, in turn, will change how often, as well as how, they use it. Finally, changing business conditions, world events, etc. will change usage patterns as well, and are totally beyond our ability to predict, let alone estimate their effects on Portal use.

Designing a complete Portal installation

Once we have the computational requirement determined, we use this to predict the number of processors needed to satisfy the expected use. In this operation, we draw upon our experience and understanding of the Portal's design to determine the best mix of systems for the customer. In general, there are two approaches:

- The "big box:" Portal and (usually) the supporting server software reside on a single departmental or "enterprise" server. This server may have as many as 32 processors, extremely fast disk storage, and a large amount of temporary storage.
- The "Portal pool": The Portal part of the installation uses a number of relatively low-cost servers. The backend database and storage system may require clustered or mirrored hardware to provide maximum availability. In addition, a load balancing unit (hardware appliance or a server with appropriate software) is used to hide which servers are in use and which ones are unavailable due to maintenance or repair, and to provide the highest availability overall for the Portal installation

Method

Rating Portal performance

Portal performance measurements require that a standard workload representative of the expected use be used, and the server(s) performance on that workload be measured over an appropriate length of time. Instantaneous performance will not be representative of real-world performance, and it may also be appropriate to measure performance after an extended period of use to determine overall fitness for purpose. While a system's performance may drop after an extended period use, significant decreases may indicate issues with the Portal software design, the operating system's ability to support extended use, or the hardware itself.

Scale out: Performance increase as servers are added

One important test is to determine whether Portal performance increases as additional hardware is added. The best result is if the performance increases linearly as servers are added. This is important because it will determine whether a server pool is an appropriate design for the Portal installation.

In the benchmark tests of the current version of Portal, it was found that overall performance changed linearly as the number of Portal servers changed from 1 to 3. This confirmed that Portal's design would scale out in a typical installation.

Scale up: Performance increase as processors are added

Similarly, the performance should increase as additional processors are added. Portal is Java-based, and makes full use of that language's support for multiple threads of execution. As additional processors are added, we would hope to see a near-linear increase in performance, which would help determine whether a single server would be appropriate for the Portal design.

Earlier benchmarks on several server platforms indicated that Portal's performance improved as additional processors were made available. The improvement was slightly non-linear, which was predicted by the scale-up limitations of the Java virtual machine used by the Portal. While the Java VM has excellent support for multiple processors in most situations, garbage collection is still limited to the use of a single processor.

This means that relatively large systems with 12 or more processors will see a drop in performance proportional to the amount of time needed for periodic garbage collection. This represents a fundamental limit on the scale up of Portal servers, unless an alternate Java VM is used with multiprocessor support for garbage collection.

Other effects

Different requests to the Portal will require different resources to satisfy. While good design can minimize, for example, the number of requests made to the Portal's backend database, certain requests will still require that the database or repository be accessed before they can be satisfied. Examples include listing the contents for a specific folder or channel, publishing a document to the Portal, and creating a new user session.

Another significant effect is the number of simultaneous requests at any time. In addition to any limits due to the Java VM or servlet container used by Portal, a drop in performance as concurrent requests increase would tend to limit the size of Portal server that can be used, which would affect the number of servers required to satisfy the predicted load.

In addition to the effects of different hardware or workplace designs, other systems can affect the Portal's performance. Clearly, for a large Portal site, the network performance, both locally and over the enterprise's Wide Area Network (WAN) can have a significant effect on performance as observed by the Portal's users. Intermediate systems, such as Web proxy servers or network routers, can also affect overall performance. Performance monitoring, both at the server and, if possible, at representative user sites, should be part of the operational plan for any enterprise Portal installation, but are well beyond the scope of this document.

Estimating user load

Once we have a rating for the Portal hardware under consideration, we need to estimate the load that users will generate. The fundamental problem here is that usage changes during the work day, and changes unpredictably based on external factors, such as business conditions, world news, or the local weather. To accommodate external and uncontrolled changes, a safety factor is chosen for the Portal's size based on the importance the Portal site has to the agency/business's operations. To accommodate the predictable portion of the Portal's load, we estimate the peak usage for the system, and use that to determine the computing power needed for the Portal installation.

"The worst time of day": Current peak usage

Analysis of Web logs from current intranet systems are usually used to point to peak usage periods in different time zones and the resultant load distribution over a day. Some of that type of thinking should be done here. (Similar analyses can be done with existing internet deployments for B2B and B2C solutions).

Session length

Session length is difficult to estimate, but some attempt should be made. Analysis of various intranet sites estimates that the typical session lasts about 10 minutes. However, it may not be representative of the Portal's use for a given application. To address these unknowns we will present several scenarios later in this document.

Template and workplace design effects

Portal design can have a significant effect on the computing power needed. Page components (e.g. static images and JavaScript files), many often invisible to the average user, have to be loaded into the user's browser to provide a complete workplace page. This effect will be small in most cases, since these components will be cached by most browsers and *because most Portal servers can deliver these components substantially faster than the more expensive workplace "portlets."*

In addition, some workplace portlets are much more "expensive" to generate than others, and may need to be used sparingly.

To accommodate this, we have performed benchmarks for two situations: an "empty workplace" without any components on it, approximating the situation where component contents are provided by other servers, and a "full workplace" with six components representing the range of tasks performed by Portal components. The latter will produce a significantly higher estimate for the computing power required for the portal solution.

Getting to a predicted equivalent load

Once we understand the effects of different components on the Portal's performance, we can sum the "cost" of each component and normalize it by the "empty workplace" performance to produce an "equivalent load" rating for a given workplace design.

Sizing the Portal installation

Safety factors

The safety factor chosen for the Portal size is, fundamentally, a function of the importance of the Portal system to the business or agency, and the confidence that the Portal use will be governed by factors that have been included in the sizing calculations. The safety factor is also important in smaller Portal installation where the failure of single server may mean the loss of substantial portion of the Portal's capacity.

User perception of Portal speed

Clearly, we could choose a smaller system if we felt that users would accept significantly slower performance during peak periods. In the current effort, slow performance could have a direct, negative effect on the Portal's use, which would go directly against the purpose of deploying the Portal system in the first place.
For this sizing study, we will ignore the effects of network latency and bandwidth. While significant for remote users, they can be minimized by careful template design. It may be appropriate, for example, to provide a "minimal" template with smaller and fewer images as an alternative for dialup or remote users.

Scope of prediction

The estimates in this study are based on benchmarks using CleverPath Portal, version 3.5_002. An enterprise-class server was used for the database backend to minimize any effects on overall performance.

Benchmarked performance is based on tests using IBM's WebSphere application server, version 4.0. Expected performance for Portal with the included Tomcat 3 servlet engine is roughly 50% of Websphere's. In addition, tests with large numbers of simultaneous requests show that Tomcat's performance is also reduced significantly with more than 100-200 simultaneous requests. Given the differences between stress testing and real user loads, this suggests that Tomcat should be "tuned" or replaced with another servlet engine if the number of concurrent sessions exceeds 500 for extended periods.

Scenarios for 10,000 user Portals

Normal use of a portal designed for 10,000 users.

Analyses of intranets estimate that the typical session length of a user session is about ten minutes.

Portal equivalent load

Assuming that content pages are comprised of static images and static (although regularly updated) text. We expect several "click throughs" for each workplace component, so we will see an equivalent load of 5-10 "requests" per workplace . The higher load would include the use of dynamic content in the components, such as channel lists.
Each workplace "view" will take the average user roughly 1 minute to read and act upon (based on previous Portal experience). As a result, a ten minute session will generate a total load equivalent to 50-100 empty workplace requests. If we have 10,000 users all accessing the portal during a peak hour, that will result in 0.5 to 1.0 million requests per hour.

Increased use: Improved interface increases demand

Increased use due to improved design

Other intranet sites have reported 100-200% increases in use as the result of improved design and new capabilities. The increases are due to additional users as well as longer user sessions. We can use the higher end of the observed increases to estimate the load as Portal use increases.

Portal equivalent load

Given a base equivalent load of 50-100, the load as use increases could reach an equivalent load of 150-300 "empty workplaces" per session, or 1.5 to 3 million request per hour.

Usage pattern changes: New system changes how intranet is used

New usage pattern

If we expect that the manner in which the Portal is used will be significantly different from a typical intranet's use, we can no longer use our existing user session length and peak period estimates. For this estimate, we will arbitrarily predict that use will increase 200% over the use expected in the second scenario.

Portal equivalent load

The Portal equivalent load will increase by 200%, and is estimated as 4.5 to 9 million requests per hour

Required hardware

Our current tests used Pentium III processor based servers (700 MHz in four processor servers). Servers used were "departmental-class" servers (Dell 64xx series).
Empty workplace performance was measured at 100 requests/second per processor or 360,000 per hour.

As a result, the number of processors required for scenario 1 will range from 1 to 3 processors (or up to 4 with a 25% safety factor on the high end of the work load).

For the second scenario – the increased use scenario – we would expect the complete system to require 4 to 8 processors, or up to 10 with the safety factor.

The final scenario – usage pattern changes – would require 12to 25 processors, or 30 or more with the safety factor.

These predictions are necessarily imprecise, as they depend on the exact intranet page design, usage rates and patterns, etc. for a large corporation. However, they confirm that current hardware should be acceptable for the first scenario, with additional hardware needed for the expanded use scenarios.

Portal installation design

Other issues

Database server requirements

Portal does require a backend database. Previous experience predicts that the database will need about 1/5 to 1/3 as many processors as are available to the main Portal. The selection and configuration of the backend database is beyond the scope of this document.

Big box approach

There are two approaches to Portal installation design. The "big box" approach reduces administrative costs (and possibly, license costs) by installing the required computing power in a single server.

The load predicted by the third and maybe the second scenarios points at the need for a UNIX-based server, since Windows 2000 servers are to be limited to 8 or sixteen processors unless the infrequently used "Datacenter Edition" is used.

Given the predicted processor requirements, the big box approach is feasible for all the scenarios. Sun, for example, sells several servers that will accommodate 8 or 24 processors (UltraSPARC III, 900 MHz).

The expected performance of the USIII should be about 1.5-2x the PIII-700MHz performance measured by our benchmark. An 8 processor V880, for example, could support the second scenario without a problem. A 24 processor Sun Fire 6800 would be needed for the third scenario. For higher load scenarios, a "big box" configuration might not be appropriate, given the cost of Sun's high-end servers.

Commodity server approach

In contrast to the big box approach, Portal can also be configured to use a pool of relatively inexpensive servers. This can increase the total software license cost and administrative costs, but will significantly reduce the hardware costs for the Portal system.

Three Dell 4xPIII servers, for example, could handle the second scenario without trouble, and, with the faster processors now available (1.5 GHz and beyond), could probably accommodate the third scenario as well.

With Sun servers, we could use low-end Sun Fire 280s with 2 900 MHz processors each. 2-3 such servers would be required for the second scenario, with as many as 10 for the third. If the third scenario is used for the planned Portal, it would be good to consider using a somewhat larger Sun server, such as the 4-processor Sun Fire 420R to reduce the total number of servers in the pool.

Recommendation Summary

The hardware estimates are listed in Table 1. This table shows the three scenarios and the two approaches to Portal site configuration that were discussed in the previous section. We've chosen to rate Sun's UltraSPARC-III processors as 50% faster than the PIII processors tested by our benchmarks. Please note that we do not recommend the use of Intel servers for the third scenario.

The conversion of a predicted number of Intel processors into a specific hardware recommendation is more of an art than a science. We really need to undertake some performance tests on typical hardware our clients use in their infrastructure. A benchmark test, such as the one that resulted in the predicted ability for a PIII 700 MHz cpu to handle 100 requests per second, covers specific, repeatable situations similar to real use. A performance test attempts to simulate "the real world" as closely as possibly A performance test requires:

- Representative amounts of data/content
- Real hardware
- Realistic workplace pages
- A defined goal: e.g., does this server support XX users without taking more than NN seconds per workplace page?

Another shortcoming in our test information is that the current benchmarks were only conducted on Intel processors. There are future plans to do similar tests on Unix boxes. In the meantime we are need Unix experts out there to help us come up with reasonable equivalences….
There are future plans to do Unix benchmarking

Table 1: Sample configurations for usage scenarios discussed earlier for 10,000 user portal

Required computing power for a 10,000 user portal		*Distributed configuration (Pool of servers)*		*Single server (Big box)*	
Scenarios	Processors	Intel	Sun	PIII-700 MHz	UltraSPARC-III 900 MHz
Normal use	2-4	2 1-way, to 2 2-way		2-4	2-3

Increased intranet use	5-10	3-5 2-way 2-3 4-way	2-4 2-way 1-2 4-way	5-10	4-7
New usage pattern (but not emergency peak)	15-30	4-8 4-way	5-10 2-way 3-5 4-way	Not recommended	10-20

Using a Sizing Calculator

A Sizing Calculator spread sheet is attached. The calculator allows one to input assumptions around number of users, % of different types of users (heavy, light, etc.), and the concurrency of those different types of users (what % are using the portal during the peak hours).

The numbers used in the attached spread sheet correspond to the "Normal use of a portal designed for 10,000 users." situation described earlier in the sizing guidelines.

Assumptions

The following sections briefly describe the assumptions that must be inserted into the sizing calculator (see the block of information in the lower left of the spread sheet on "model description").

World Wide Portal User Population

This is the total number of users the client expects to use the portal deployment.

Distribution of Types of Users Within the Total Population

Users are categorized as "very light", "light", "average", and "heavy" users. This is where the first major assumption is made. The load profile for these four types is accounted for in the next two assumptions.

Rough Concurrency Guide

This is the first of the load profile assumptions. It is a very rough assumption of what percent of a given type of user is logged on to the portal and having an active session during the peak hour of the day. For an average intranet user, as an example, probably half of the users in a given time zone log on to the portal and have an active session during the first official work hour of the day. The distribution of users in a given time zone over the work day may be as shown in Table 2. (This is based on some rough guesses based on experience: intranet use is heaviest at the start, mid-morning, and around lunch, with a small peak in the afternoon. This would be for a generic intranet, not a project management portal or something more specialized.) This table indicates that if the population is split evenly in each time zone, then the peak concurrency would only be 32.5%. This is the kind of analysis that needs to be thought out to arrive at the concurrency numbers to be inserted into the sizing calculator. For the 10,000 user example discussed above the very simple assumption was that all users were on during the same hour.

Table 2: Concurrency of users during the day

Time of day in time zone	Fraction of users on portal during the hour
7 am	20
8	50
9	30
10	20
11	30
12 noon	40
1 pm	30
2	20
3	20
4	30
5	20
6	10

Table 3: % of total user population on portal during each hour of day[1]

Time of day in first time zone	Fraction of users on portal from all 4 time zones during most of the day	Sum of fractions	% of total users
8 am	50+20+0+0	70%	
9 am	30+50+20+0	100%	25%
10 am	20+30+50+20	120%	30%
11 am	30+20+30+50	130%	32.5%
Noon	40+30+20+30	120%	30%
1 pm	30+40+30+20	120%	30%
2 pm	20+30+40+30	120%	30%
3 pm	20+20+30+40	110%	27.5%
4 pm	30+20+20+30	100%	
5 pm	20+30+20+20	100%	
6 pm	10+20+30+20	80%	
7 pm	00+10+20+30	60%	
8 pm	00+00+10+20	30%	

[1] Assume that the population is split evenly between the time zones… if the assumption is different, just apply the appropriate math to account for that distribution.

Portal Equivalent Load Factors for each type of user

This is the last major assumption about the user load profile. This captures the amount of portal activity, expressed in PEL/second or requests per second, that each user type is involved in during a typical hour. Some of the activity that affects these numbers are

- The amount of time the user is actively using the portal (length of the session time during the hour)
- How many clickthroughs, or how many links a user clicks on per workplace
- How many workplaces a user looks at and acts on during the session

In the scenarios above on "Normal", "Increased usage due to improved design of intranet", and "Usage changes: new system changes how intranet is used", it was described how major changes in the way the intranet is used could affect loading patterns.

In the first of these scenarios it was described how one could expect several "clickthroughs" for each workplace component, so one will see an equivalent load of 5-10 "requests" per workplace. The higher load would include the use of dynamic content in the components, such as channel lists. Furthermore it was assumed that each workplace "view" will take the average user roughly 1 minute to read and act upon (based on previous Portal experience). As a result, an average user who has one "typical" ten minute session during the hour, will generate a total load equivalent to 50-100 empty workplace requests or PEL per hour.

In the sizing calculator the assumptions for the three user types were chosen as shown:

- Heavy user: 200 PEL/hour
- Average user: 100 PEL/hour
- Light user: 50 PEL/hour
- Very light user: 25 PEL/hour

Again, these assumptions have to be explored for the client situation at hand. As described in the scenarios above the average user numbers doubled for each of the scenarios as usage increased. The average user PEL factor of 100 generates roughly the same result as described in the "normal" scenario above. Changing that PEL factor to 300 and 900 in the calculator for average users for the increased usage scenarios will result in predicting roughly the same number of cpu's as described in the scenarios above.

Conclusion

The sizing calculator has several assumptions that need to be considered for the client scenarios at hand. This sizing guideline hopefully provides guidance on how to consider these assumptions.

The bad news is that performance or benchmark tests have not been conducted on Unix machines, nor have they been conducted with commercial application servers other than WebSphere.

The good news is that the CleverPath Portal SWAT team in development is embarking on performance tests for a number of leading application servers, using Portal 4.0. This will help us to make some better sizing estimates.

Sizing Calculator

Sizing calculator-replaces calculator in TechTalk 15 May 02

(last update: 3 July 02, see 3 July 02 sizing guide by Lynch, Melstner, and Annis for guidance)

total user poplulation ww = 10000

1. Distribution of types of users within ww population

	B2E	B2B	B2C
Heavy	0.0%		
Average	100.0%		
Light	0%		
Very Light	0%		

2. Rough Concurrency Guide

	B2E	B2B	B2C
Heavy	0%		
Average	100%		
Light	0%		
Very Light	0%		

3. Number of Concurrent Users

	B2E	B2B	B2C
Heavy	0		
Average	10,000		
Light	0		
Very Light	0		

4. Portal Equivalent Load Factors per Hour

	B2E	B2B	B2C
Heavy	200	20	5
Average	100	5	1
Light	50	1	0.20
Very Light	25	0.20	0.04

Benchmark performance

Quad P-III 700 MHz = 400 PEL's per second			
	Quad Rating	2800	power units
	Standard	**0.142857143**	PEL/sec/unit

5. Total PEL's per second

	B2E	B2B	B2C
Heavy	0	0	0
Average	278	0	0
Light	0	0	0
Very Light	0	0	0
Total load	**278**		

6. Intel Machines (MHz) (when using a commercial application server… see Sizing Model description on lower left)

CPU's	400	500	600	700	800	900	1000	1250	1500
1	4.86	3.89	3.24	2.78	2.43	2.16	1.94	1.56	1.30
2	2.43	1.94	1.62	1.39	1.22	1.08	0.97	0.78	0.65
3	1.62	1.30	1.08	0.93	0.81	0.72	0.65	0.52	0.43
4	1.22	0.97	0.81	0.69	0.61	0.54	0.49	0.39	0.32
6	0.81	0.65	0.54	0.46	0.41	0.36	0.32	0.26	0.22
8	0.61	0.49	0.41	0.35	0.30	0.27	0.24	0.19	0.16
10	0.49	0.39	0.32	0.28	0.24	0.22	0.19	0.16	0.13
12	0.41	0.32	0.27	0.23	0.20	0.18	0.16	0.13	0.11
13	0.37	0.30	0.25	0.21	0.19	0.17	0.15	0.12	0.10
14	0.35	0.28	0.23	0.20	0.17	0.15	0.14	0.11	0.09
15	0.32	0.26	0.22	0.19	0.16	0.14	0.13	0.10	0.09
16	0.30	0.24	0.20	0.17	0.15	0.14	0.12	0.10	0.08
17	0.29	0.23	0.19	0.16	0.14	0.13	0.11	0.09	0.08
18	0.27	0.22	0.18	0.15	0.14	0.12	0.11	0.09	0.07
19	0.26	0.20	0.17	0.15	0.13	0.11	0.10	0.08	0.07
20	0.24	0.19	0.16	0.14	0.12	0.11	0.10	0.08	0.06
21	0.23	0.19	0.15	0.13	0.12	0.10	0.09	0.07	0.06
22	0.22	0.18	0.15	0.13	0.11	0.10	0.09	0.07	0.06
23	0.21	0.17	0.14	0.12	0.11	0.09	0.08	0.07	0.06

Key:
- (magenta) = insufficient power
- (yellow) = possible, depends on safety factor comfort
- (green) = plenty of power (safety factor greater than or equal to 100%)

Recommendation (depending on saftey factor)

Description of Sizing Model (see Sizing Guide document of 3 July 2002):

(J1) = Assume ww population

(table 1) = Distribution of types of users within ww population

(table 2) Concurrency of user types (% that are using the portal during the peak hour)

(table 3) = # of concurrent users per type (tbl 1 * tbl 2 * J1)

(table 4) = PEL factors for user types (avg user uses portal for 10 minutes, 5-10 clicks or requests per workplace, and each workplace takes about a minute to look over, for a total of 50-100 requests in the 10 minutes the user is actively using the portal during that peak hour… this number could vary greatly… as an example, if the user is active for 20 minutes rather than 10 minutes, there will be 100-200 requests during the peak hour… also if the user quickly goes thru a workplace (say 1/2 minute rather than 1 minute) this will double the PEL

(I14) = Doculab benchmark (load handled per power unit) power unit measured in MHz of cpu

(table 5) = PEL's per second per type (tbl 3 * tbl 4)/3600

(table 6) = (tbl 5 total load) / ((I14 load/power unit)*(power unit or MHz))

The predicted sizing is based on using a commercial application server. If the default Tomcat servlet engine is used, the number of cpu's should be increased as much as 100% under certain conditions…

Basis for usage assumptions (see Sizing Guide document of 3 July 2002):