



e-business

Best Practices for High Volume Web Sites

December 12, 2001

Linda Legregni
legregni@us.ibm.com
Program Director, High Volume
Web Sites Project Office
IBM Software Group

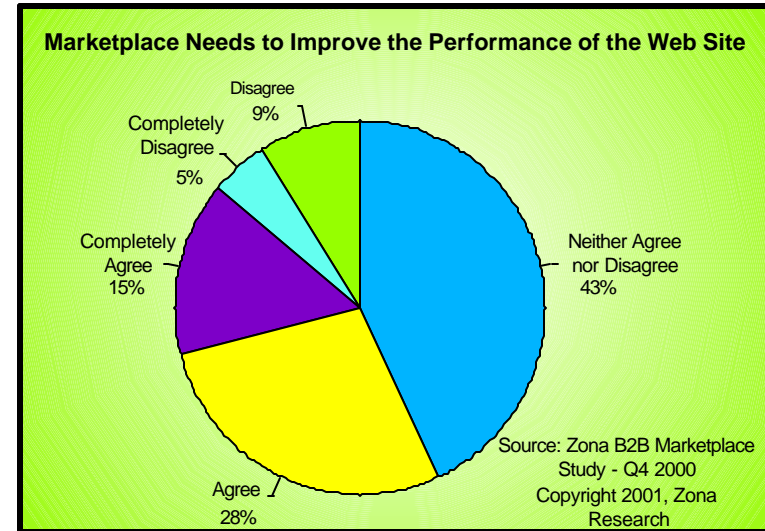
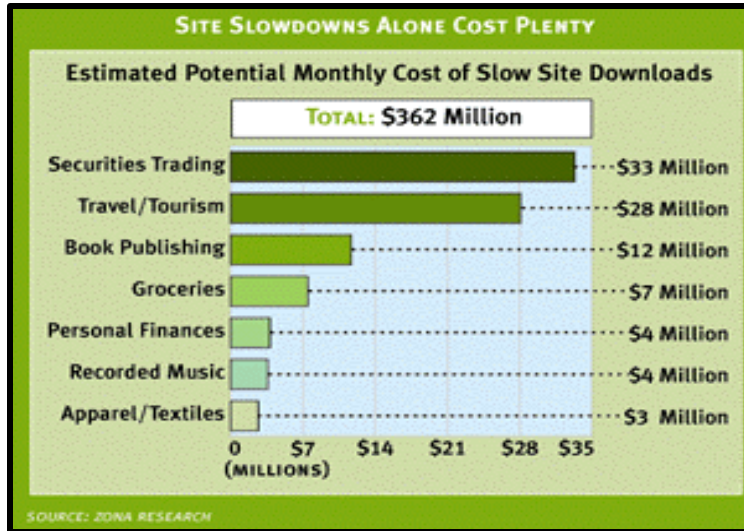




e-business

Slow Performance Costs \$\$

Zona Research Studies



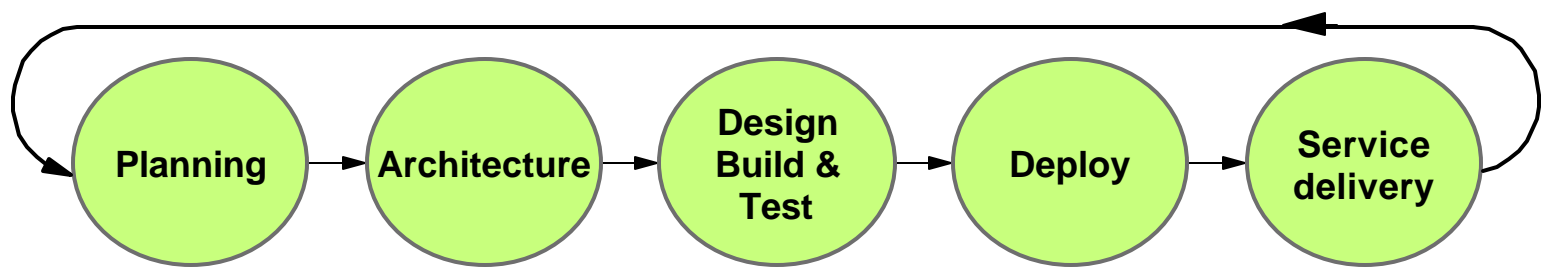
- Slow performance costs e-commerce sites \$362 million per month.
- Average Web Buyer will wait 8 seconds for a page to download. Longer download times can cause 30% of online buyers to bail out.
- One-day outages could cost \$2.8 million for an auction site, and \$5.2 million for a Net brokerage.
- I/T survey (4Q2000) : 43% says web site needs performance improvement.





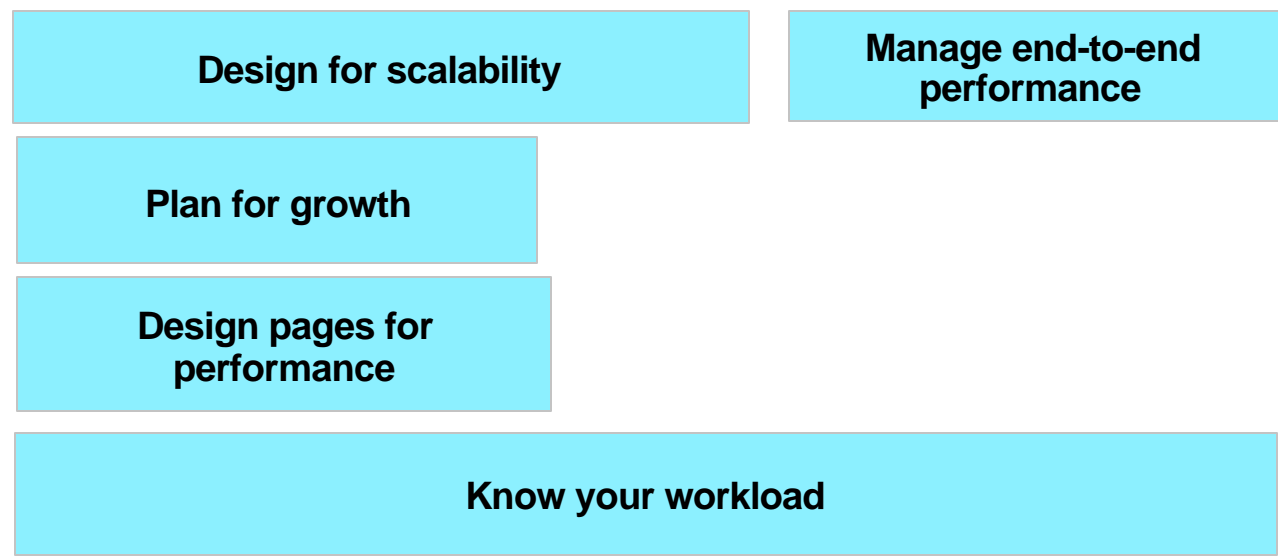
e-business

High Volume Web Sites Best Practices



Life Cycle of a Web Site

Best Practices



Note: Material regarding the topics addressed in this presentation may be obtained at:
"ibm.com/websphere/developer/zones/hvws"

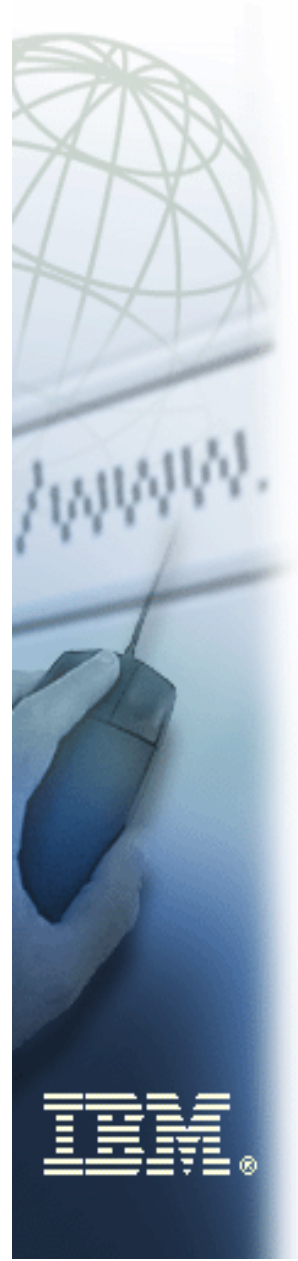




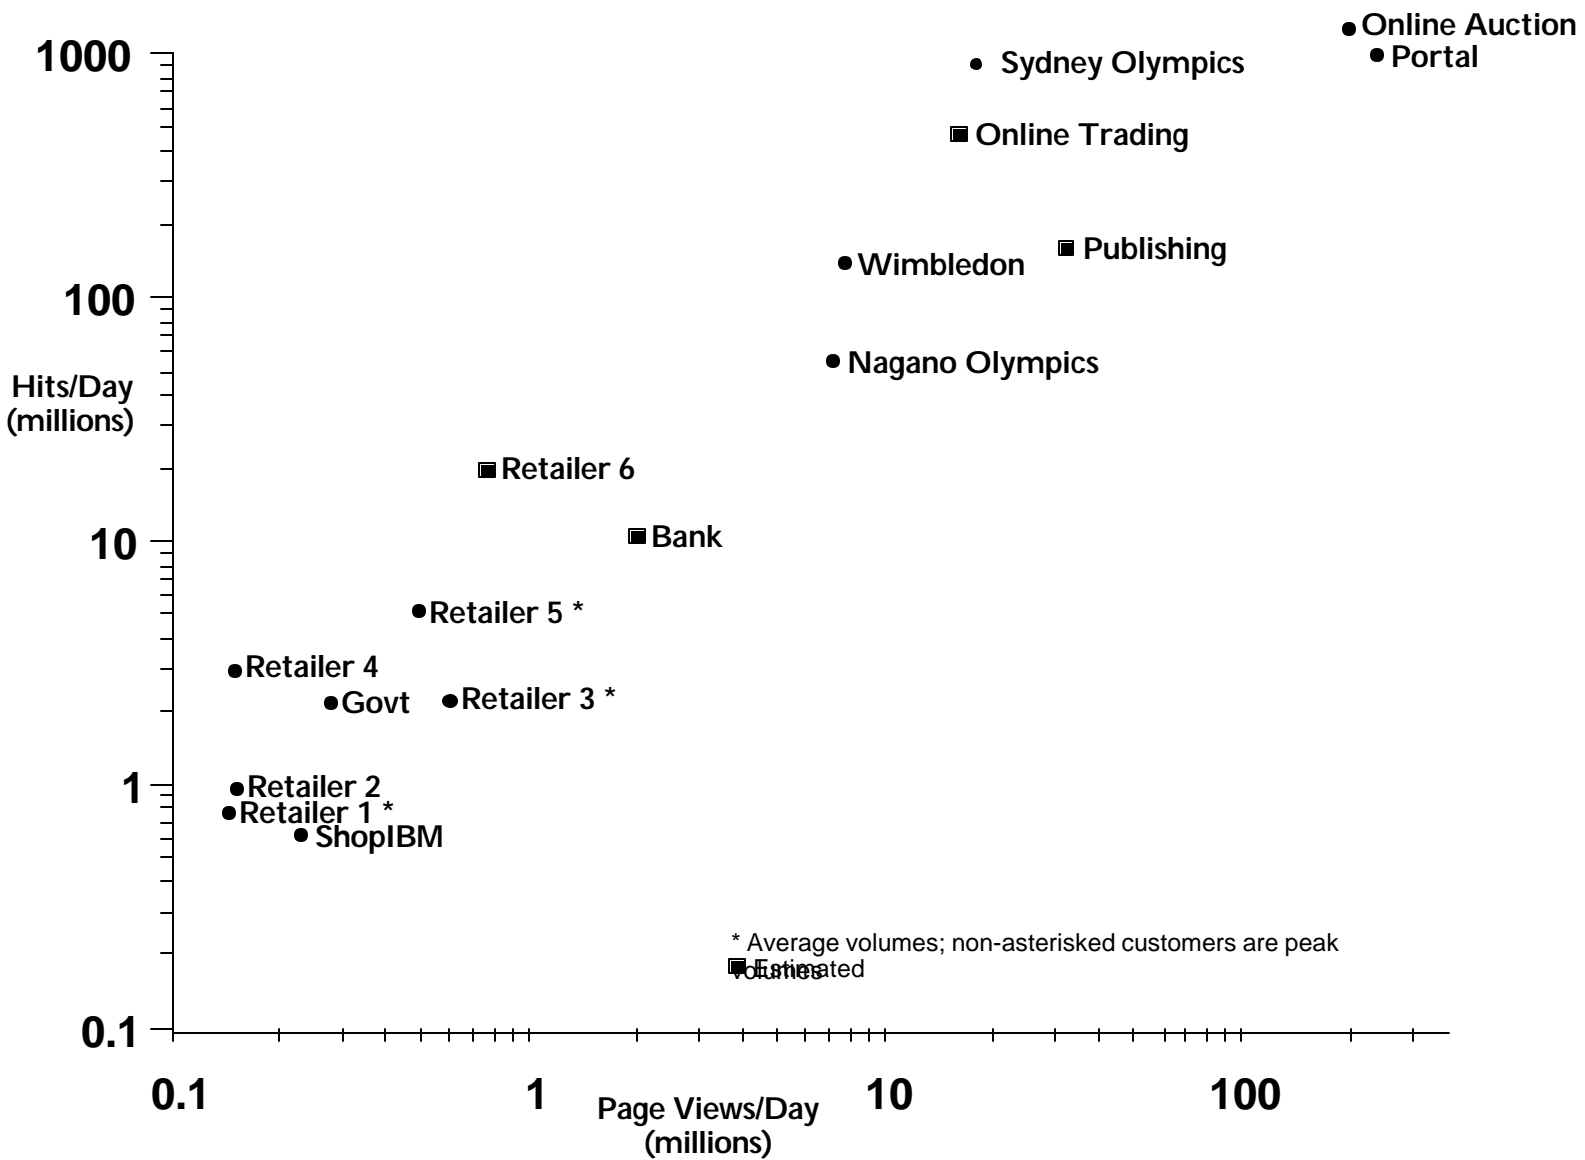
e-business

Patterns for e-Business & Workload Taxonomy

<p>Publish & Subscribe (User-to-Data)</p>	<p>Customer Self-Service (User-to-Business)</p>	<p>Online Trading (User-to-Business)</p>
<p>Online Shopping (User-to-Online Buying)</p>	<p>Business-to-Business (Business-to-Business)</p>	



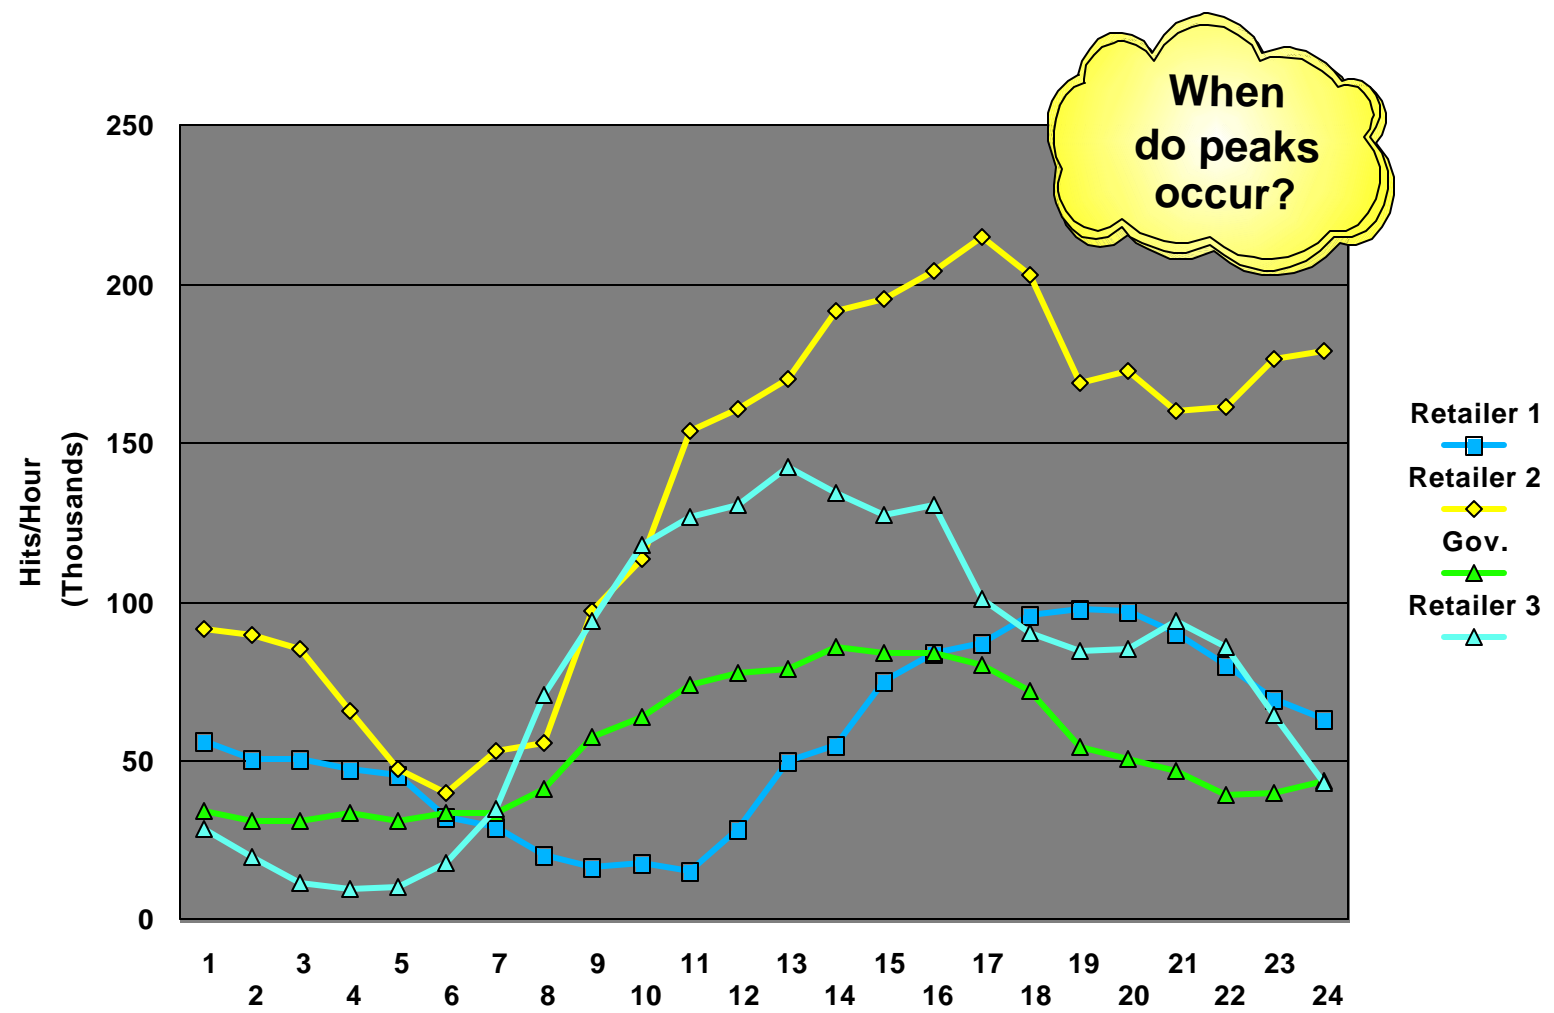
Customer Metrics Examples Hits & Page Views Per Day



Web Site Workloads Usage Pattern Over One Day



e-business

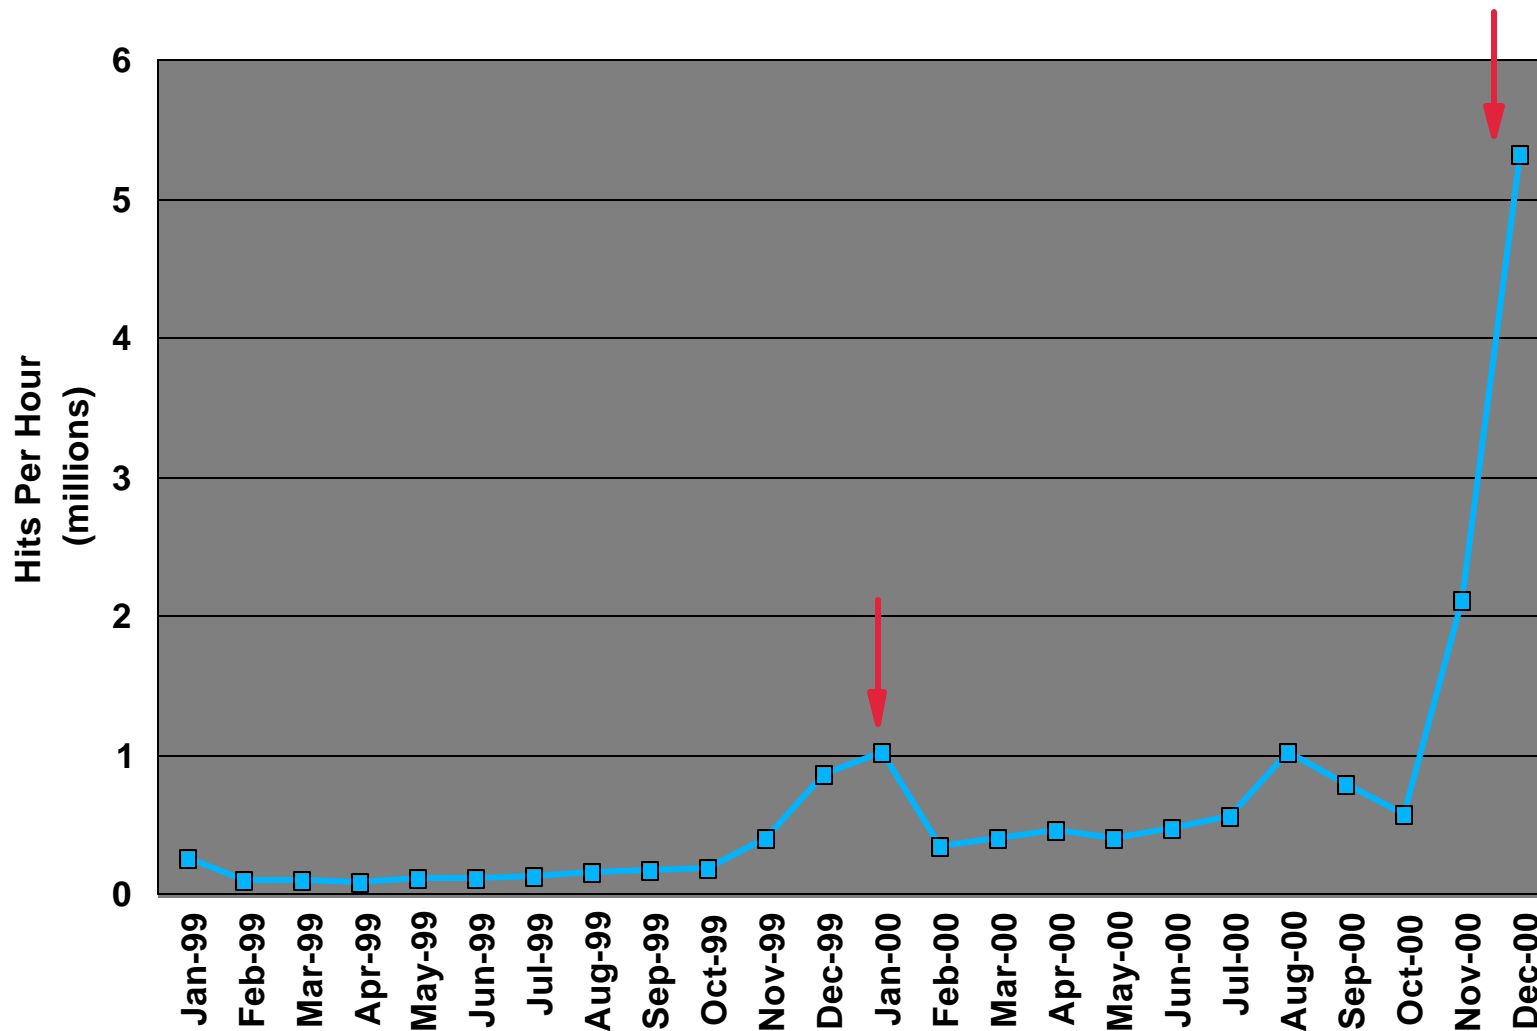




e-business

Web Site Workloads

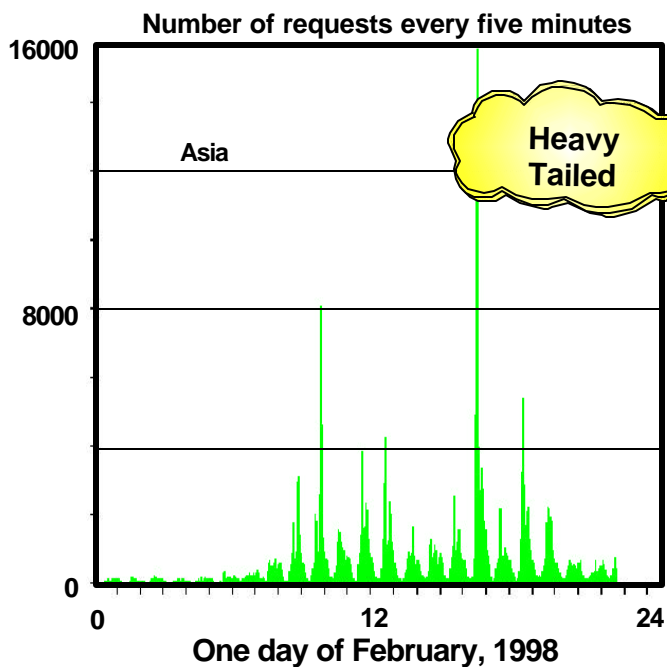
Seasonal Effects at an Online Retailer





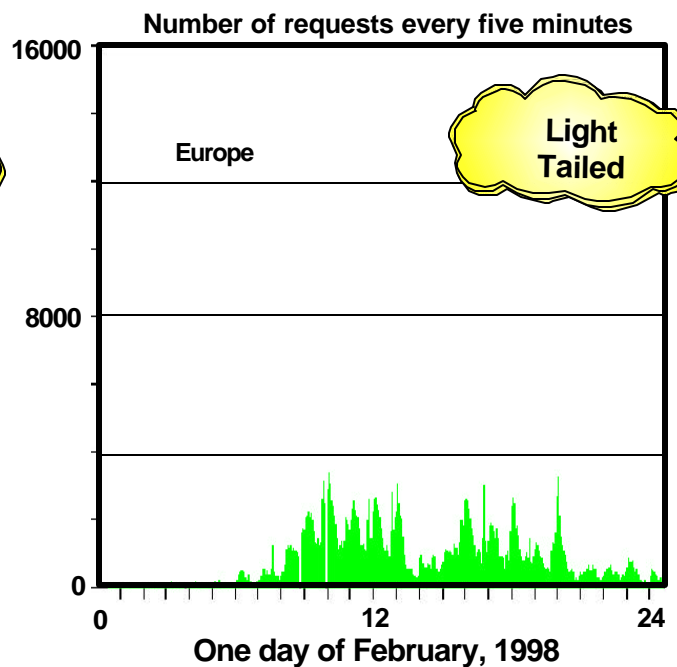
e-business

Understand Web Site Workloads Bursty Traffic Patterns 1998 Nagano Olympic Games



Contains very large bursts

Large traffic bursts - strong interest in Japan for events related to ski jumping



Much less bursty

Traffic more scattered due to time differences and broader interests in events

Sydney Olympics much less bursty due to 18 hour broadcast delay





e-business

Design for Scalability Best Practices

Six steps to scaling:

- Understand the environment
- Categorize your workload
- Determine the components most impacted
- Select the scaling techniques to apply
- Apply the techniques
- Reevaluate



Workload Characteristics for Scaling



System Workload Characteristics	Publish/Subscribe	Online Shopping	Customer Self-Service	Trading	Business-To-Business
Transaction Volumes	○	◐	◐	●	◐
Dynamic Content	◐	◐	◐	●	◐
Dynamic Searches	○	●	○	○	◐
User Specific Responses (Personalization)	○	○	◐	●	◐
Cross-session Information	○	●	●	●	●
Transaction Complexity	○	◐	●	●	●
Legacy Integration	○	◐	●	●	●
Transaction Volume Swings	○	◐	◐	●	●
Data Volatility	○	○	○	●	◐
Number of Content Publishers/Sources	●	○	○	●	◐
Page Content Volatility	●	●	◐	●	◐
Number of Unique Items per page	●	◐	○	◐	◐
Number of Page Views	●	◐	○	◐	◐
Percentage of Secure Pages	○	◐	◐	●	●
Security, Authentication etc.	○	●	●	●	●

○ Low ◐ Medium ● High

Refer to "Design for Scalability" White Paper for further information
["ibm.com/websphere/developer/zones/hvws"](http://ibm.com/websphere/developer/zones/hvws)





e-business

How Do Workload Characteristics Impact Solution Components?

System Workload Characteristics	Web Servers	Web Appl Servers	Network	Security Servers	Fire-walls	Ent. Inf. Systems	Data Base Server
Transaction volumes	○	○	○	◐	○	●	●
Dynamic content	●	◐	◐	◐	○	○	○
Dynamic searches	●	●	◐	○	◐	◐	●
User specific responses (personalization)	○	●	○	○	○	●	●
Cross-session information	◐	●	○	○	○	○	◐
Transaction complexity	○	●	○	◐	○	●	●
Legacy integration	○	●	◐	◐	○	●	●
Transaction volume swings	◐	●	○	○	○	●	●
Data volatility	●	●	○	○	○	◐	●
Number of content publishers/sources	●	○	●	○	○	○	○
Page content volatility	●	●	◐	○	○	○	●
Number of unique items per page	○	◐	○	○	○	●	●
Number of page views	●	○	●	○	●	○	○
Percentage of secure pages	●	○	○	●	○	○	○
Security, authentication, etc.	●	●	○	●	○	●	○

○ Low

◐ Medium

● High

Refer to "Design for Scalability" White Paper for further information
"ibm.com/websphere/developer/zones/hvws"

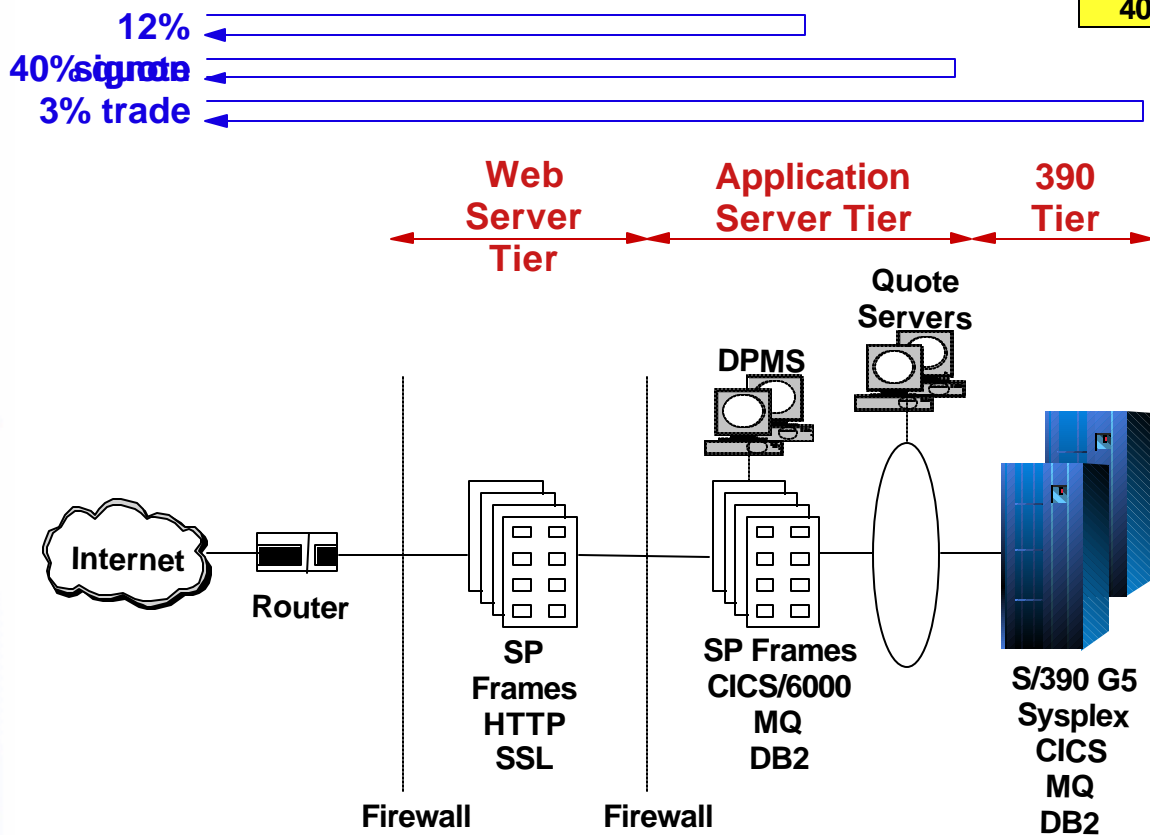


Highly Scalable Online Stock Trading Web Site



e-business

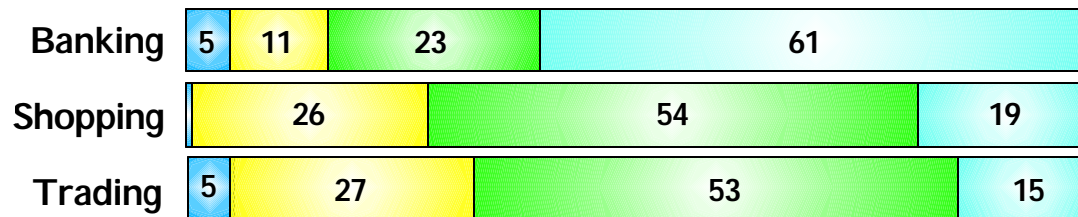
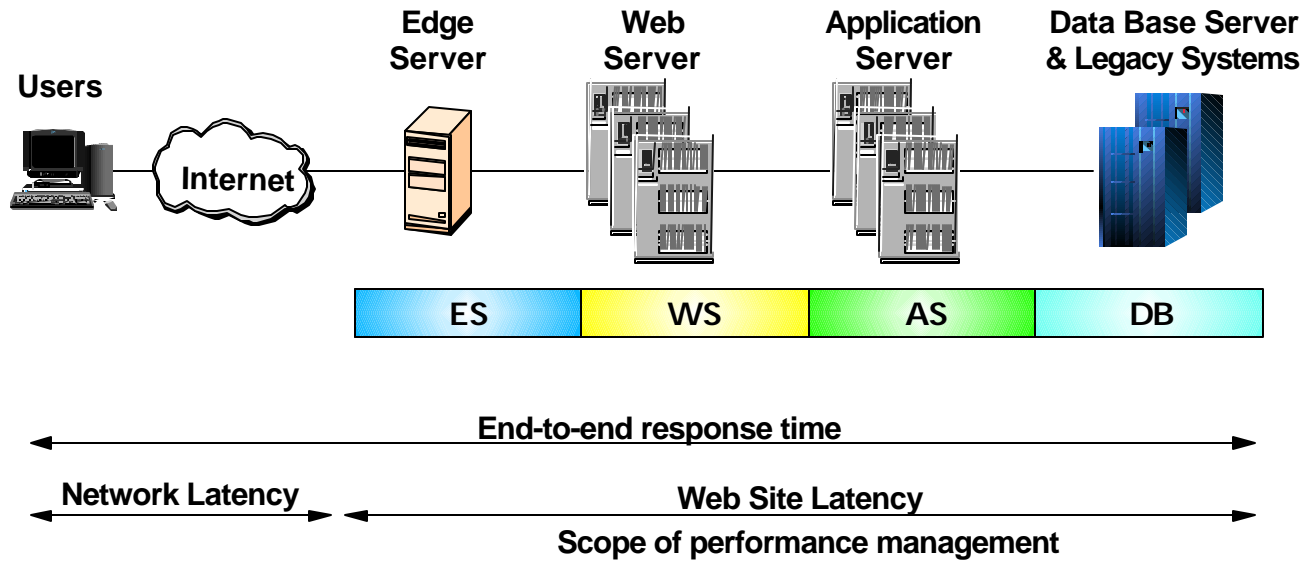
1/00	Secure site
110K	Peak concurrent users
500K	Peak trades/day
10.7M	Peak transactions/day
80M	Peak hits/day
4000	Peak trans/sec





e-business

Where do you focus your scalability and tuning?



% Latency Examples

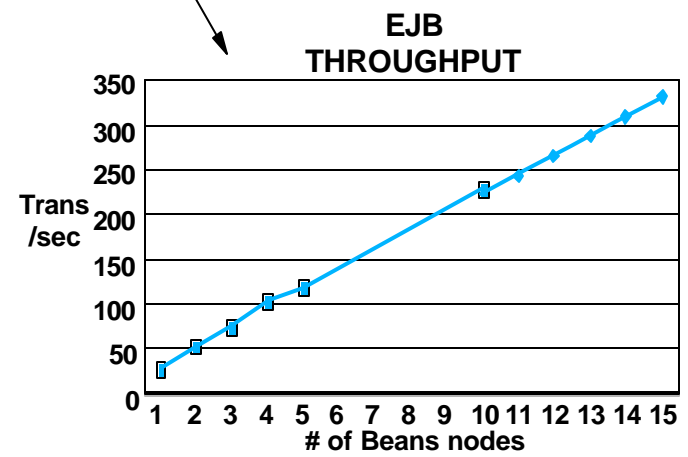
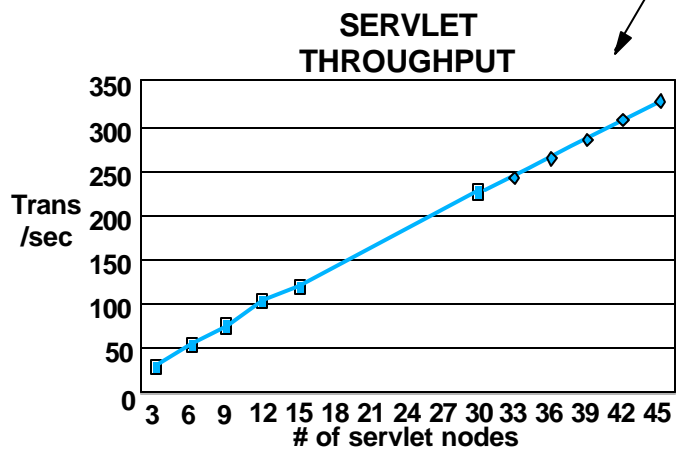
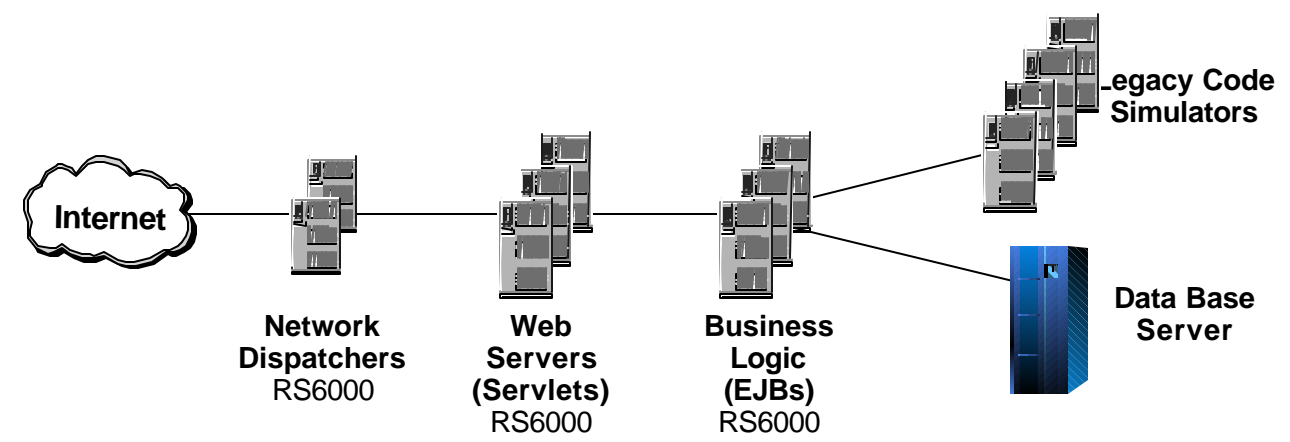
- Edge Server
- Web Server
- Application Server
- Data Base Server





e-business

Scaling Characteristics of EJBs and Servlets in an Online Trading Environment



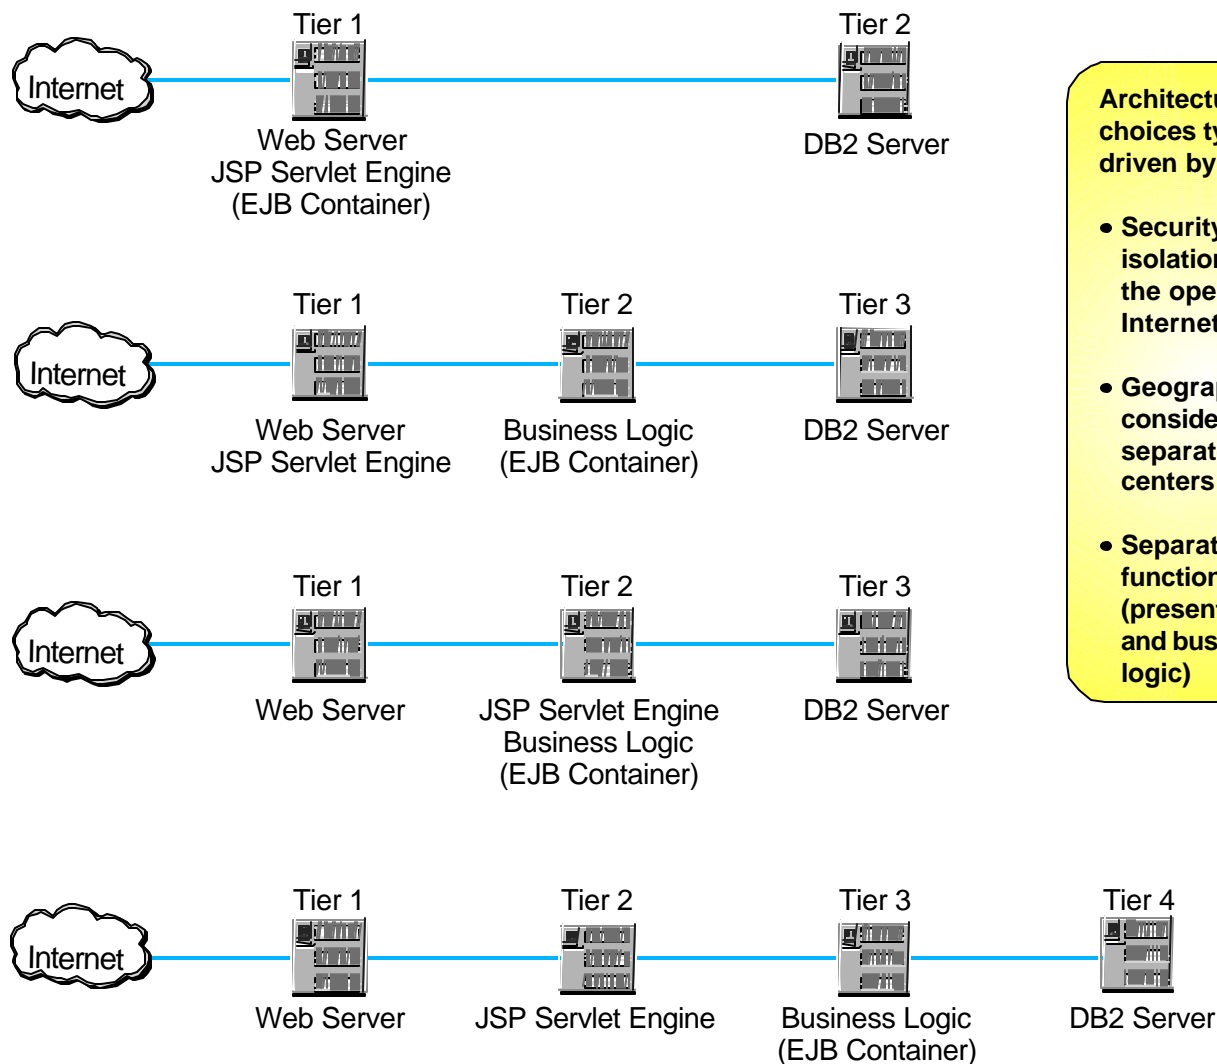
- Up to 3x faster than CGI
- Scaling is linear
- Web to EJB server ratios
- # of servers per domain
- operations procedures





e-business

WebSphere Architecture Alternatives



Architectural choices typically driven by:

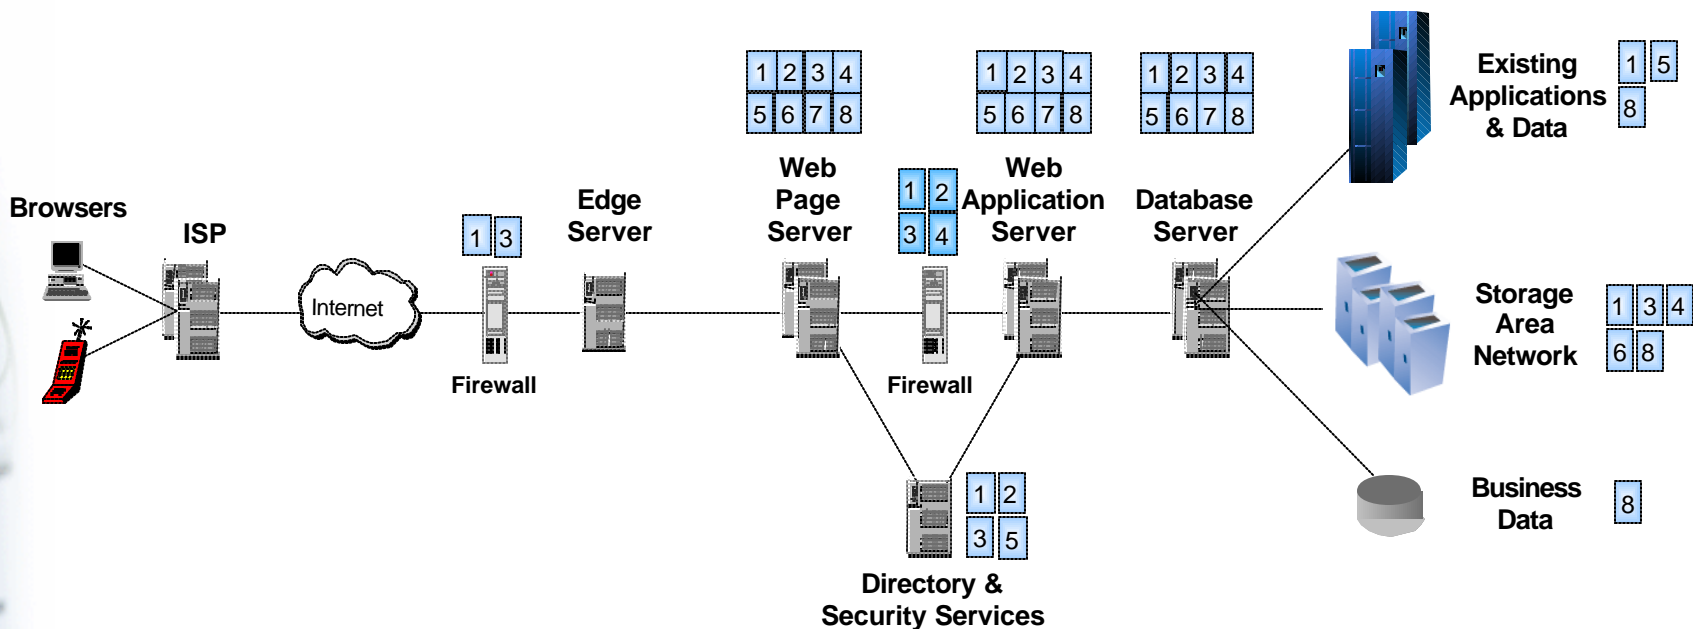
- Security, isolation from the open Internet
- Geographical considerations - separate data centers
- Separation of function (presentation and business logic)





e-business

Design for Scalability Best Practices



Techniques

- 1 Faster Machine
- 2 Replicate Machines
- 3 Specialized Machines
- 4 Segmented Workload
- 5 Request Batching
- 6 User Data Aggregation
- 7 Connection Management
- 8 Caching

Best practice: Follow a methodology to optimize for scalability

- Understand thoroughly your application environment
- Categorize your workload
- Determine the components most impacted
- Select the scaling techniques to apply
- Apply the techniques



e-business

Planning for Growth Best Practices

- Determine your workload characteristics
 - Trends
 - Burstiness
- Plan your infrastructure from end to end
 - Leads to balanced solutions
- Integrate capacity planning into your I/T and business processes





e-business

HVWS Simulator

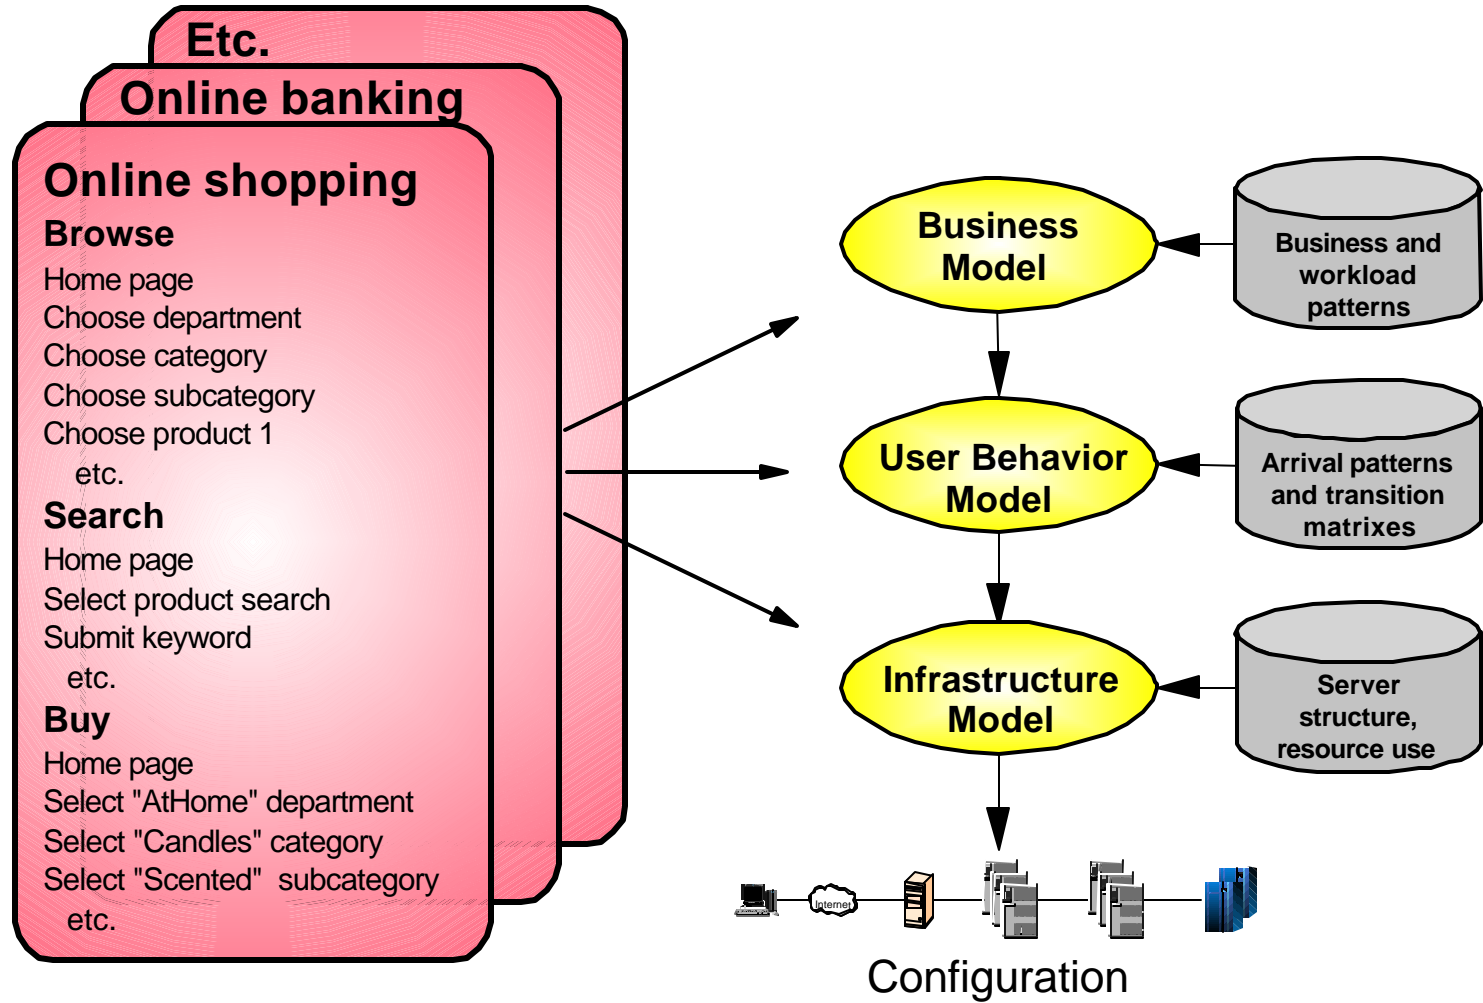
- Model to estimate performance & capacity of Web site infrastructure
 - Analyze and recommend architecture
 - "What if" analysis of performance & capacity
 - Driven by business pattern workloads
 - WebSphere/DB2 focus
 - Special algorithms to handle burstiness of Web traffic
- Available for use by IBM sales and service





e-business

Capacity Planning Process with the HVWS Simulator





e-business

Web Page Design Best Practices

Average server response time	<0.5 sec
Number of items per page	<20
Page load time	<30 sec *
Page size in bytes	<64K

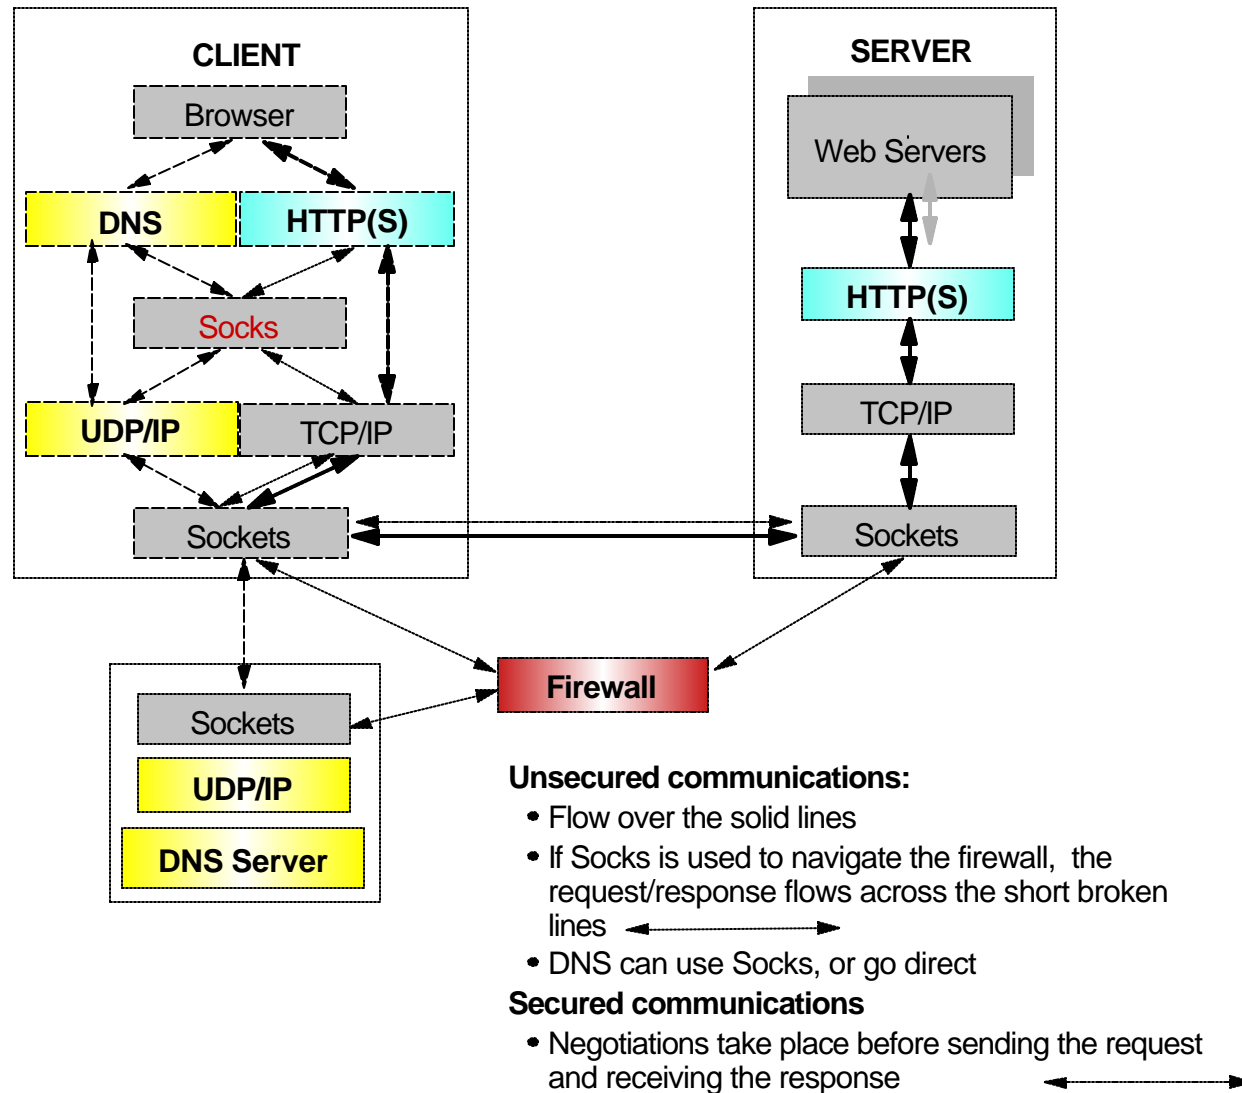
Seconds to load *	Ranking
Less than 10	Excellent
10-15	Very good
15-20	Good
20-25	Adequate
25-30	Slow
Over 30	Unacceptable

* Based on 56K connection speed



e-business

Browser & Web Server Interactions





e-business

IRS Home Page

- Very slow page load time
- Menu should be displayed early so users can navigate quickly
- Long page forces scrolling -- consider use of horizontal white space

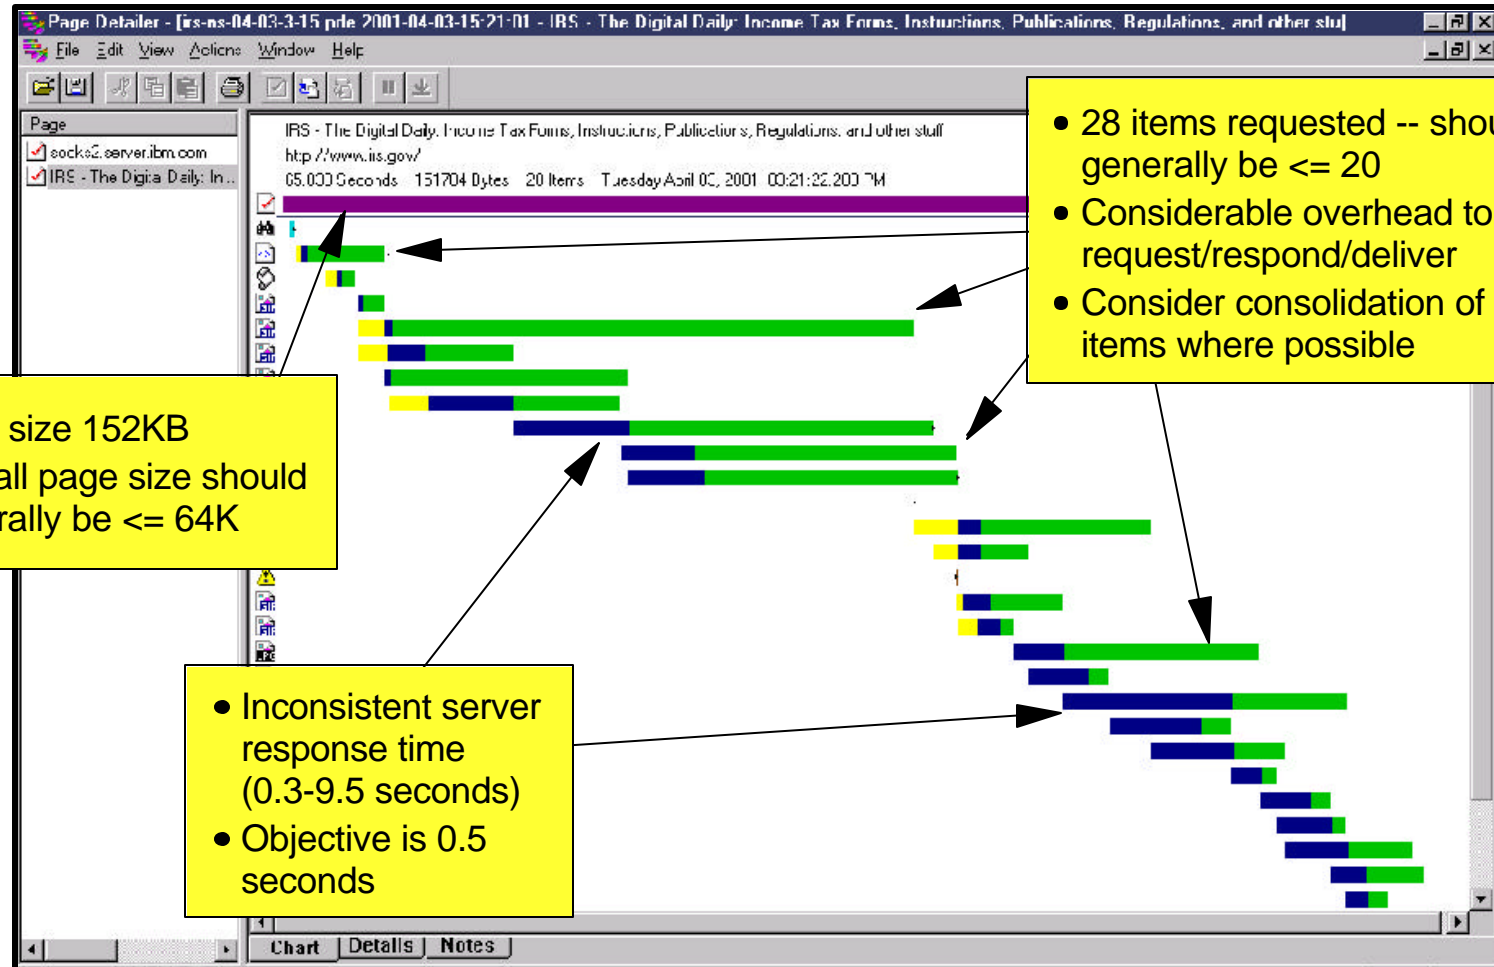
	4/30/00
Download Time	68.8 sec
Page Size	152 KB
Page Items	28





e-business

IRS Page Detailer Analysis





e-business

WebSphere Studio Page Detailer Analysis

The screenshot shows the Page Detailer window for the URL `http://www.irs.gov/`. The main area displays a waterfall chart with colored bars representing different resources. The 'Item Properties' dialog is open, showing the following details:

Item Properties		
Summary Sizes Headers Timing Events Comments		
Item	<code>http://www.irs.gov/bin/cd-si..</code>	
Content-Type	<code>image/gif</code>	
13264	Bytes	
12.977	Total Time	
4.498	Server Response Time	
8.479	Delivery Time	
328	Socket ID	
192.239.92.40:80	Remote Address:Port	
32.227.173.191:1083	Local Address:Port	

At the bottom right of the Page Detailer window, a summary table shows the following data:

28	157166	37.549
----	--------	--------





e-business

Summary of Key Measurements

	Page load time (sec)	Page size (bytes)	Number of items	Number of connections	Number of servers	Number of failed connections
Retailer 1	32.33	179,968	51	17	2	0
Retailer 2	30.5	140,842	80	7	2	0
Govt.	68.84	151,704	28	9	1	2
Retailer 3	26.26	122,146	53	7	1	0
Retailer 4	78.26	121,664	56	21	3	0
Retailer 5	41.648	111,281	37	5	2	0
News	34.45	105,433	35	21	2	0
Retailer 6	22.18	93,580	29	6	1	0
Retailer 7	22.52	84,240	46	46	1	0
Finance	27.03	72,411	36	36	4	0
Bank	19.951	64,347	30	19	1	0
Bank 2	29.741	61,073	40	11	1	0
Bank 3	15.14	56,430	25	5	1	0
Bank 4	15.69	43,891	23	23	1	0
Ebroker	8.77	39,189	12	5	2	0





e-business

WebSphere Studio Page Detailer - Availability

- Part of WebSphere Studio Advanced 3.5



Design Pages for Performance Best Practices



e-business

- **Page items**
 - Minimize number, size, and complexity of page items -- consider the business value of each
 - Store and retrieve items used more than once from the browser cache
- **Servers**
 - Access as few as possible -- avoid redirection
 - Keep connections open -- use persistent connections
 - Use the same level of code across servers
- **Avoid unnecessary use of**
 - SSL -- Assign private information to private pages and secure private data only
 - Banner ads
 - Dynamic pages
- **Use preproduction utilities that remove extra white space from the source HTML**
 - E.g., GZIP
- **Structure sequences to fully exploit parallelism**
 - Request items early





e-business

HVWS White Papers

www.ibm.com/websphere/developer/zones/hvws/

- **Design for Scalability - An Update**
 - **Describes component selection and management techniques you can use to make your Web site ready to adapt to increasing traffic**
- **Planning for Growth**
 - **Introduces a methodology for capacity planning, identifying workload patterns and configuring site infrastructure**
- **Design Pages for Performance**
 - **Discusses what drives Web page download times**
 - **Introduces page design practices that can reduce download time and improve resource utilization**
- **Managing Web Site Performance**
 - **Discusses end to end performance management**
 - **Self management**
- **Web Site Personalization**
 - **Introduces current and future techniques for personalizing your Web site**
- **Schwab puts growth plan to the test**
 - **Describes a joint IBM/Schwab project to develop an architecture that can cope with soaring growth**

