Session 214552573
City of Sacramento
PeopleSoft 8.9 Infrastructure Planning
Wednesday, April 16th 9:45am

Presenter Name: Ray Songco
Title: Technical Director
CherryRoad Technologies Inc.

Agenda

- Project Scope
- Sizing Methodology
- Architecture Design
- Q & A

About the City

- The capital of the State of California and the county seat of Sacramento County.
- Seventh most populous city in California with a 2007 estimated population of 467,343.
- The core cultural and economic center of its four-county metropolitan area (El Dorado, Placer, Sacramento, and Yolo counties) with a combined population of 2,103,956.
- The Sacramento Metropolitan Area is the largest in the Central Valley, and is the fourth-largest in California, behind the Greater Los Angeles Area, the San Francisco Bay Area, and the San Diego area.
- Greater Sacramento has been cited as one of the five "most livable" regions in America, and the city was cited by Time magazine as America's most integrated.

Project Scope

Electronic Citywide Accounting and Personnel System (eCAPS)

- Financials Phase I: Financials 8.9, Portal 8.9
- HCM Phase I: HCM 8.9
- Financials Phase II: Budgeting, Cash/Deal Management, Vendor Self-Service
- HCM Phase II: ELM, Employee Self-Service, Candidate Gateway, Talent Acquisition Manager, eApplications

Project Scope

Other Applications (non-eCAPS)

- Revenue Management – CIS
- Customer Relationship Management – Siebel
- Land Management – Accela
- Identity Management – TBD
- Enterprise Job Scheduling – Tidal

Project Scope

Infrastructure Planning

- Before
  - IT Standards
  - Hardware Sizing
  - Redundancy
- During
  - Load Testing
  - Security
- After
  - Production Environment Administration
  - Multi-Phase Strategy

Project Scope

IT Standards

- Platform
  - SUSE Linux/Oracle 10g
  - HP Proliant Series
- High Availability
  - Multiple Domains
  - Content Switching
  - Disaster Recovery
- Security
  - LDAP Integration
  - Multiple Firewalls
  - SSL Encryption

Sizing Methodology

Sizing Approach

Step 1: Estimate Peak Utilization (PSTPM)

Step 2: Calculate Database Server Utilization

Step 3: Estimate Database Server Memory

Step 4: Calculate Application Server Utilization

Step 5: Estimate Application Server Memory

Step 6: Calculate Web Server Utilization

Step 7: Estimate Web Server Memory

Step 8: Determine Network Bandwidth Requirements

Sizing Methodology

Step 1: Gather Peak Utilization and Calculate Peak Online PeopleSoft Transactions per Minute (PSTPM):

- Heavy
- Medium
- Light
- Self-Service

Estimate Peak/Peak Utilization	*Heavy*	*Medium*	*Light*	*Self-Service**	*Total*
PSTPM	2	1	0.5	0.1	
Peak Online	76	-	276	74	426.13
Peak Online PSTPM	152	-	138	7	**297.40**

Sizing Methodology

Step 2: Calculate Database Server Utilization

- Add 20% for daytime batch processes
- Calculate TPC-C Requirement, Total Peak PSTPM

Database Server Utilization					
Add 20% Batch Overhead:					
Total Peak PSTPM					**356.88**
Required TPC-C Rating @ 250 TPC-C/PSTPM					
Peak TPC-C Requirement					**89,218.80**
Reference Server Benchmark:					
(HP Integrity rx5670/Oracle10g/RHL AS3) TPC-C					136,110.00
Reference vs. Actual Server SPEC_Int Differential					5.63%
Estimate Actual TPC-C Rating from Differential					**143,778.17**
(HP DL580G3/Oracle10g/SUSE 9) Est. TPC-C					
Peak DL580G3 CPU Utilization					**62.05%**

Sizing Methodology

Database Server

- Look up www.tpc.org for TPC-C rating

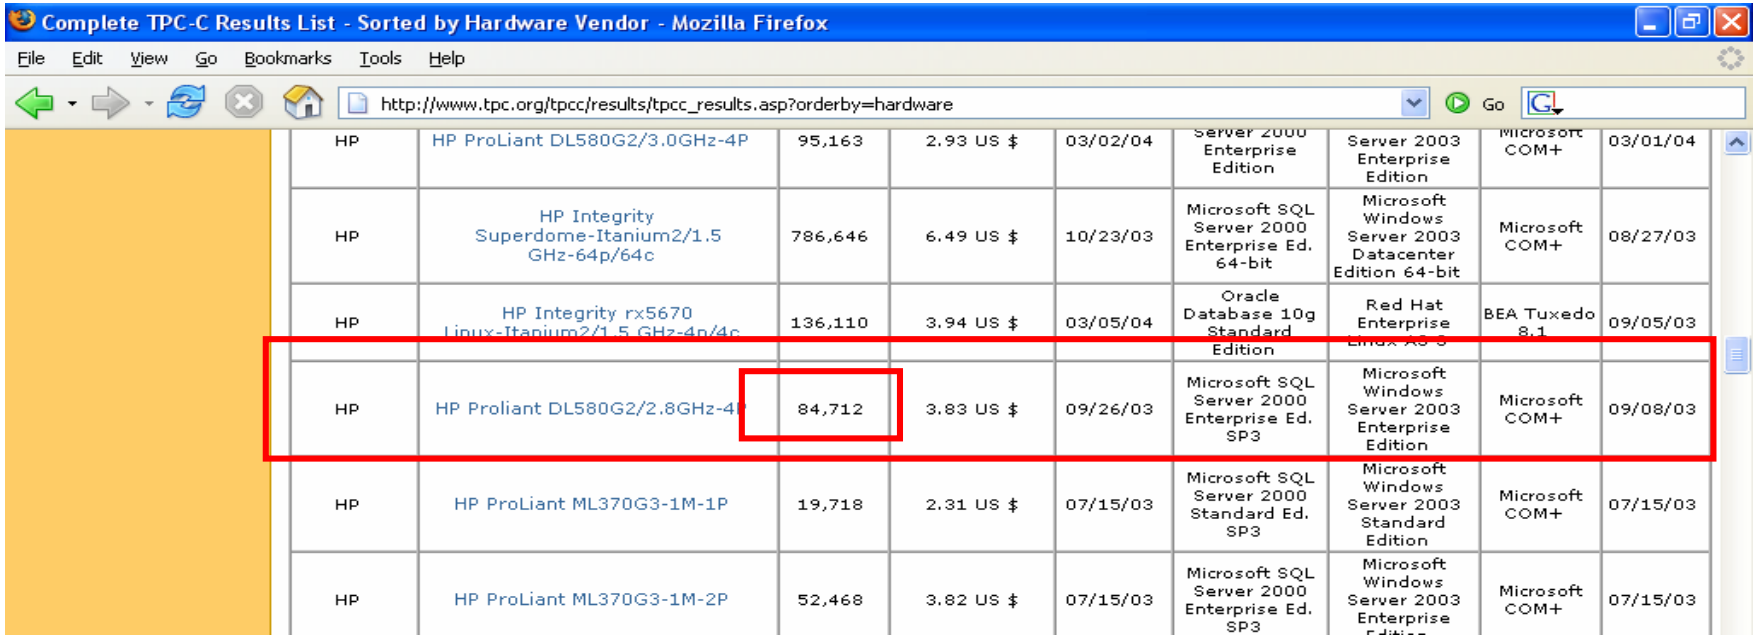

HP	HP ProLiant DL580G2/3.0GHz-4P	95,163	2.93 US $	03/02/04	Server 2000 Enterprise Edition	Server 2003 Enterprise Edition	Microsoft COM+	03/01/04
HP	HP Integrity Superdome-Itanium2/1.5 GHz-64p/64c	786,646	6.49 US $	10/23/03	Microsoft SQL Server 2000 Enterprise Ed. 64-bit	Microsoft Windows Server 2003 Datacenter Edition 64-bit	Microsoft COM+	08/27/03
HP	HP Integrity rx5670 Linux-Itanium2/1.5 GHz-4p/4c	136,110	3.94 US $	03/05/04	Oracle Database 10g Standard Edition	Red Hat Enterprise Linux AS 3	BEA Tuxedo 8.1	09/05/03
HP	HP Proliant DL580G2/2.8GHz-4P	84,712	3.83 US $	09/26/03	Microsoft SQL Server 2000 Enterprise Ed. SP3	Microsoft Windows Server 2003 Enterprise Edition	Microsoft COM+	09/08/03
HP	HP ProLiant ML370G3-1M-1P	19,718	2.31 US $	07/15/03	Microsoft SQL Server 2000 Standard Ed. SP3	Microsoft Windows Server 2003 Standard Edition	Microsoft COM+	07/15/03
HP	HP ProLiant ML370G3-1M-2P	52,468	3.82 US $	07/15/03	Microsoft SQL Server 2000 Enterprise Ed. SP3	Microsoft Windows Server 2003 Enterprise	Microsoft COM+	07/15/03

Sizing Methodology

Database Server

If no exact TPC-C match found, pick a proxy and use www. spec.org to determine % differential:

67	HP Integrity Server rx4640 (1500 MHz, Itanium 2)	4 cores, 4 chips, 1 core/chip	63.4	64.2
68	HP Integrity Server rx5670 (1300 MHz, Itanium 2)	1 core, 1 chip, 1 core/chip	12.2	12.4
69	HP Integrity Server rx5670 (1300 MHz, Itanium 2)	2 cores, 2 chips, 1 core/chip	24.2	24.5
70	HP Integrity Server rx5670 (1300 MHz, Itanium 2)	4 cores, 4 chips, 1 core/chip	48	48.6
71	HP Integrity Server rx5670 (1500 MHz, Itanium 2)	1 core, 1 chip, 1 core/chip	15.2	15.2
72	HP Integrity Server rx5670 (1500 MHz, Itanium 2)	2 cores, 2 chips, 1 core/chip	30.3	30.3
73	HP Integrity Server rx5670 (1500 MHz, Itanium 2)	4 cores, 4 chips, 1 core/chip	60	60
74	HP Integrity Superdome (1.6GHz/24MB Dual Core Itanium 2)	128 cores, 64 chips, 2 cores/chip (Hy	2367	2367
75	HP Integrity Superdome (1.6GHz/9MB Itanium 2, 16 cells)	64 cores, 64 chips, 1 core/chip	1108	1108
76	HP Integrity Superdome 16-way (1500 MHz Itanium 2)	16 cores, 16 chips, 1 core/chip	229	229
77	HP Integrity Superdome 32-way (1500 MHz Itanium 2)	32 cores, 32 chips, 1 core/chip	453	453
78	HP Integrity Superdome 64-way (1500 MHz Itanium 2)	64 cores, 64 chips, 1 core/chip	904	904

	From spec.org				
Server	Model	#CPU	MHz	SPEC_RATE	SPEC_RATE/MHz/CPU
Database	HP Integrity rx5670 (Reference)	4	1500	60	0.0100000
Database	HP DL580G3 (Actual)	2	3000	56.8	0.0094667
	Difference	**5.63%**			

Sizing Methodology

Database Server

- Extrapolate TPC-C Rating from Differential
- Calculate Peak CPU Utilization

Required TPC-C Rating @ 250 TPC-C/PSTPM					
Peak TPC-C Requirement					**89,218.80**
Reference Server Benchmark:					
(HP Integrity rx5670/Oracle10g/RHL AS3) TPC-C					136,110.00
Reference vs. Actual Server SPEC_Int Differential					5.63%
Estimate Actual TPC-C Rating from Differential					**143,778.17**
(HP DL580G3/Oracle10g/SUSE 9) Est. TPC-C					
Peak DL580G3 CPU Utilization					**62.05%**

Sizing Methodology

Step 3: Estimate Database Server Memory

- Estimated based on previous customer experience
- Stratified ranges based on estimated database size
  - Headcount
  - Applications

Estimate Database Memory					
Database Size	RAM (GB)				
Small 20 - 80 GB	**4 - 12**				**12**
Medium 80 - 140 GB	12 - 24				
Large 140 - 220 GB	24 - 36				

Sizing Methodology

Step 4: Calculate Application Server Utilization

- Get SPEC_Int Rating of Application Server
  - If no exact match found use proxy (see Database Server Sizing)
- Regression: 0.25 SPEC_Int Rating (SIR) per PSTPM
- Calculate Peak Utilization at SIR Requirement / SIR Capacity

Application Server Utilization					
# CPU's					4
SPEC_Int Rating					84.20
Peak Utilization @ 0.25 SIR/PSTPM					74.35
Peak DL380G4 CPU Utilization					**88%**

Sizing Methodology

Step 5: Estimate Application Server Memory

- Regression: RAM requirement based on application (e.g., HCM User = 100 MB RAM)
- Calculate Memory Requirement - round up to 2 GB increments.

GB Memory @ 0.01 GB RAM/user					
Peak Memory Requirement					4

Sizing Methodology

Step 6: Calculate Web Server Utilization

- Basic PeopleSoft guidelines set ratio of 1:4 for Application Server to Web Server

Web Server Utilization					
Required CPU @ 25% of Application Server					18.59
Peak DL380G4 CPU Utilization					22%

Sizing Methodology

Web Server

- Regression: each concurrent user consumes 5 MB RAM
- Calculate Memory based on Peak Concurrent Users

GB Memory @ 200 users / 1 GB RAM					
Peak Memory Requirement					2

Sizing Methodology

Network Capacity

- PIA Servers - Dual Gigabit
- Internal Clients – 100 MBPS
- Self-Service – Sufficient ISP Bandwidth:

Bandwidth	End Users
56Kb	8-10
64Kb	10-12
128Kb	20-24
256Kb	40-48
512Kb	80-96
T1 (1.544Mb)	~240
T3 (44.736Mb)	~7,000

Sizing Methodology

Total Capacity

- Aggregate metrics for each Application Suite
- If hardware already procured, calculate peak utilization levels:

Component	Peak Utilization		
Database Server	**FMS/Portal**	**HCM**	**EPM**
CPU Utilization	44.18%	62.05%	45.04%
Memory Requirement	12	12	12
Application/Web Server			
CPU Utilization	80%	110%	80.11%
Memory Requirement	12	6	8

Sizing Methodology

Total Capacity

- For Infrastructure Design, provide hardware options based on selected vendor:

Combined Totals:	
Database Server:	
1)	HCM SIR + CRM SIR = 200 +12.35 = **212.35 SIR**
	Sun Fire E6900 (16 processor) 32 cores, 16 chips, 2 cores/chip = **204**
	Sun Fire E2900 (12 processor) 24 cores, 12 chips, 2 cores/chip = **255**
	Sun Fire E4900 (12 processor) 24 cores, 12 chips, 2 cores/chip = **257**
	Sun Fire E6900 (16 processor) 32 cores, 16 chips, 2 cores/chip = **230**
	Sun Fire E6900 (16 processor) 32 cores, 16 chips, 2 cores/chip = **249**
Application Server:	
	1) HCM SIR + CRM SIR = 262.5 + 16.2 = **278.7 SIR**
	2) Divide by number of servers in cluster:
	278.7 SIR / 2 = **139.35 SIR**
	Sun Fire V40z 8 cores, 4 chips, 2 cores/chip = **136**
	Sun Fire E4900 (8 processor) 16 cores, 8 chips, 2 cores/chip = **173**
	Sun Fire E6900 (12 processor) 24 cores, 12 chips, 2 cores/chip = **189**
Web Server:	
	1) HCM SIR + CRM SIR = 65.625 + 4.05 = **69.675 SIR**
	3) Divide by number of servers in cluster:
	69.675 SIR / 2 = **34.8375 SIR**
	Sun Fire 3800 8 cores, 8 chips, 1 core/chip = **34.5**
	Sun Fire V20z 2 cores, 2 chips, 1 core/chip = **40.4**
	Sun Fire V440 (1600MHz) 4 cores, 4 chips, 1 core/chip = **38.7**
	Sun Fire V880 8 cores, 8 chips, 1 core/chip = **36.2**

Sizing Methodology

Final Configuration

Server Breakdown	
Database Servers:	12
Web/App Servers:	12
Windows Batch:	11
File Server:	1
Total Servers:	**36**

PRODUCTION ENVIRONMENT CITY GARAGE

PRODUCTION DATABASE SERVERS:
3 X DL580 G4 4 X 3.4GHZ DC 16MB L3 XEON 7140
2 X 72 15K RPM SAS
16GB RAM
SUSE 9 R3

PRODUCTION WEB/APP SERVERS:
3 X DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72 15K RPM SAS
12GB RAM
SUSE 9 R3

PRODUCTION CRYSTAL/BATCH SERVERS:
3 X DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72 15K RPM SAS
12GB RAM
WINDOWS 2003 R2 64-BIT

TRAINING ENVIROMENT

TRAINING DATABASE SERVER:
DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72GB R1
12GB RAM
SUSE 9 R3

TRAINING WEB/APP SERVER:
DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72GB R1
12GB RAM
SUSE 9 R3

TEST/UAT ENVIRONMENT

FINANCIALS TEST DATABASE SERVER:
DL580 G4 4 X 3.4GHZ DC 16MB L3 XEON 7140
2 X 72 15K RPM SAS
16GB RAM
SUSE 9 R3

FINANCIALS TEST WEB/APP SERVER:
DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72 15K RPM SAS
12GB RAM
SUSE 9 R3

FINANCIALS TEST CRYSTAL/BATCH SERVER:
DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72 15K RPM SAS
12GB RAM
WINDOWS 2003 R2 64-BIT

FINANCIALS TEST DATABASE SERVER:
DL580 G4 4 X 3.4GHZ DC 16MB L3 XEON 7140
2 X 72 15K RPM SAS
16GB RAM
SUSE 9 R3

HR TEST WEB/APP SERVER:
DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72 15K RPM SAS
12GB RAM
SUSE 9 R3

HR TEST CRYSTAL/BATCH SERVER:
DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72 15K RPM SAS
12GB RAM
WINDOWS 2003 R2 64-BIT

DR ENVIRONMENT COMM CENTER

DR DATABASE SERVERS:
3 X DL580 G4 4 X 3.4GHZ DC 16MB L3 XEON 7140
2 X 72 15K RPM SAS
16GB RAM
SUSE 9 R3

DR WEB/APP SERVERS:
3 X DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72 15K RPM SAS
12GB RAM
SUSE 9 R3

DR CRYSTAL/BATCH SERVERS:
3 X DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72 15K RPM SAS
12GB RAM
WINDOWS 2003 R2 64-BIT

PS_HOME FILE SERVER.
DL380 G4 2 X 2.4 DC XEON
12GB RAM
WINDOWS 2003 R2 64-BIT

FCS IMPLEMENTATION

FINANCIALS DATABASE SERVER:
DL380 G4 2 X 3.4GHZ DC
2 X 146GB 15K R1, 4 X 146GB 15K R10
12GB RAM
SUSE 9 R3

FINANCIALS WEB/APP SERVER:
DL380 G4 2 X 2.4 DC XEON
2 X 72 R1, 3 X 146 R5
12GB RAM
SUSE 9 R3

FINANCIALS CRYSTAL/BATCH SERVER:
DL380 G4 2 X 2.4 DC XEON
2 X 72 R1, 3 X 146 R5
12GB RAM
WINDOWS 2003 R2 64-BIT

HCM IMPLEMENTATION

HR DATABASE SERVER:
DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72GB R1
12GB RAM
SUSE 9 R3

HR WEB/APP SERVER:
DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72GB R1
12GB RAM
SUSE 9 R3

HR CRYSTAL/BATCH SERVER:
DL380 G4 2 X 2.4 DC XEON
2 X 72GB R1, 3 X 146GB R5
12GB RAM
WINDOWS 2003 R2 64-BIT

EPM IMPLEMENTATION

EPM DATABASE SERVER:
DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72GB R1
12GB RAM
SUSE 9 R3

EMP WEB/APP SERVER:
DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72GB R1
12GB RAM
SUSE 9 R3

EPM CRYSTAL/BATCH SERVER:
DL360 G5 2 X 1.8GHZ QUAD CORE XEON 5140
2 X 72GB R1
12GB RAM
WINDOWS 2003 R2 64-BIT

Sizing Methodology

Lessons Learned

- Be conservative! Better to oversize than undersize.
- Hardware is driven by vendor models, not metrics.
- Go to vendor online stores and “shop” for servers to gain familiarity with servers.
- Remember that estimates are exactly that, and are proven out later during load testing.
- TPC and SPEC_Int are becoming less reliable. May need to switch to other measure such as Summation of GHz.
- Non-production hardware sizing usually based on previous sizing information for similarly-scaled customers.

Architecture Design

What is Redundancy?

- **Failover** – capability to switch over automatically to a redundant or standby system upon the failure or abnormal termination of the previously system. Failover happens without human intervention and without warning (vs. Switchover).

- **Load Balancing** - load balancing is a technique (usually performed by load balancers) to spread work between many computers, processes, disks or other resources in order to get optimal resource utilization and decrease computing time.

Architecture Design

Levels of Redundancy

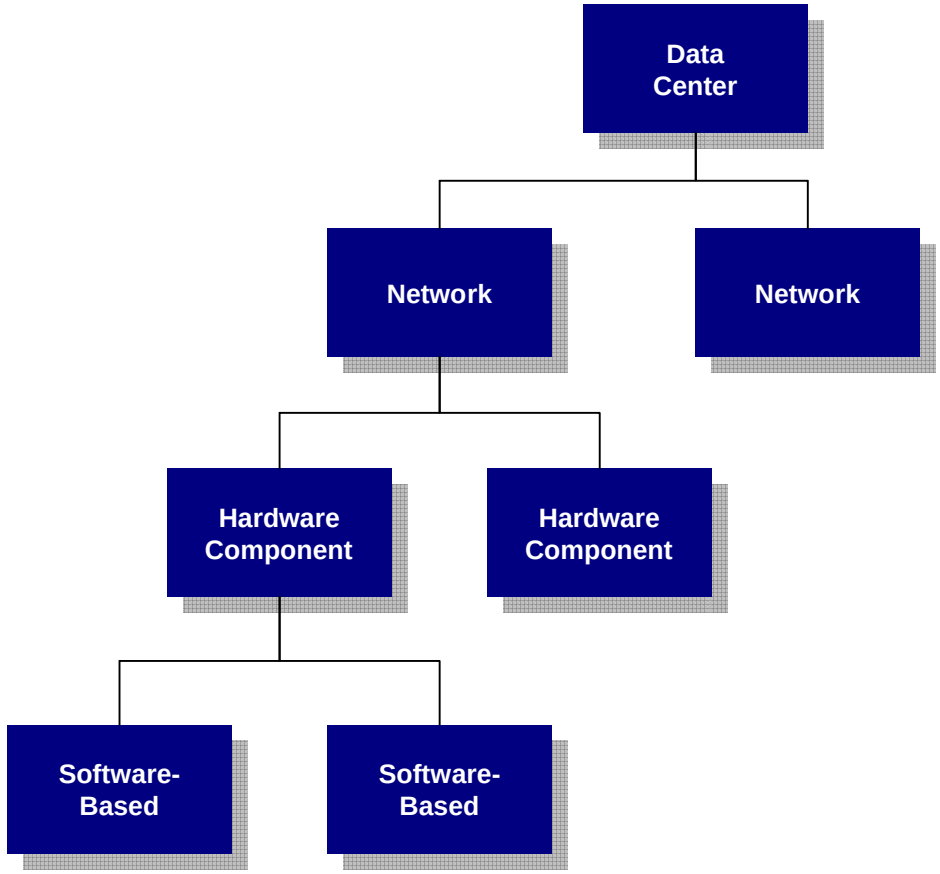

Remote Disaster Recovery site

Load Balancing/ Failover appliances, Server farms

Hardware redundancy (multiple NIC's, UPS, etc.)

Built into middleware (multiple web servers, Domains, etc.)

Architecture Design

Typical PeopleSoft Redundancy

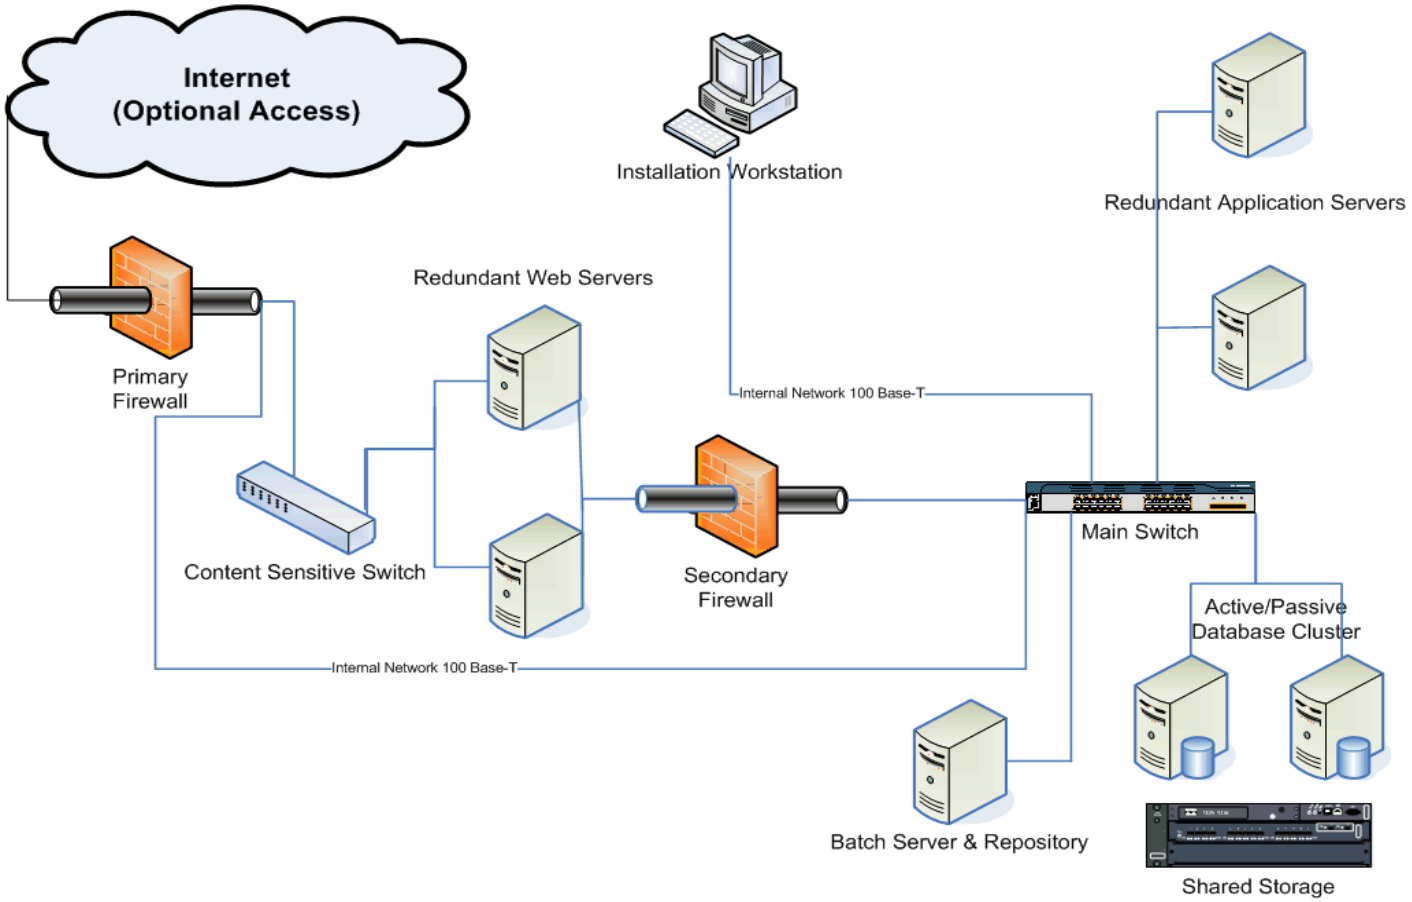

Architecture Design

Types of Scaling

- Horizontal Scaling – Multiple smaller servers grouped together by a load balancer
  - Pros – Cheaper boxes, scalability, no obsolescence, ease of administration
  - Cons – Support costs
- Vertical Scaling – Single enterprise-class server, upgradeable CPU, memory chips
  - Pros – Lower support cost
  - Cons – Expensive, parts phased out, not as scalable

Architecture Design

PIA Load Balancing/Failover

- Content Switches
  - Internal User Redundancy
  - External Enterprise Redundancy
- Disaster Recovery Site
  - Phase I: Web and Application Tiers
  - Phase II: Database Tier (RAC)

Architecture Design

Proposed Enterprise Architecture

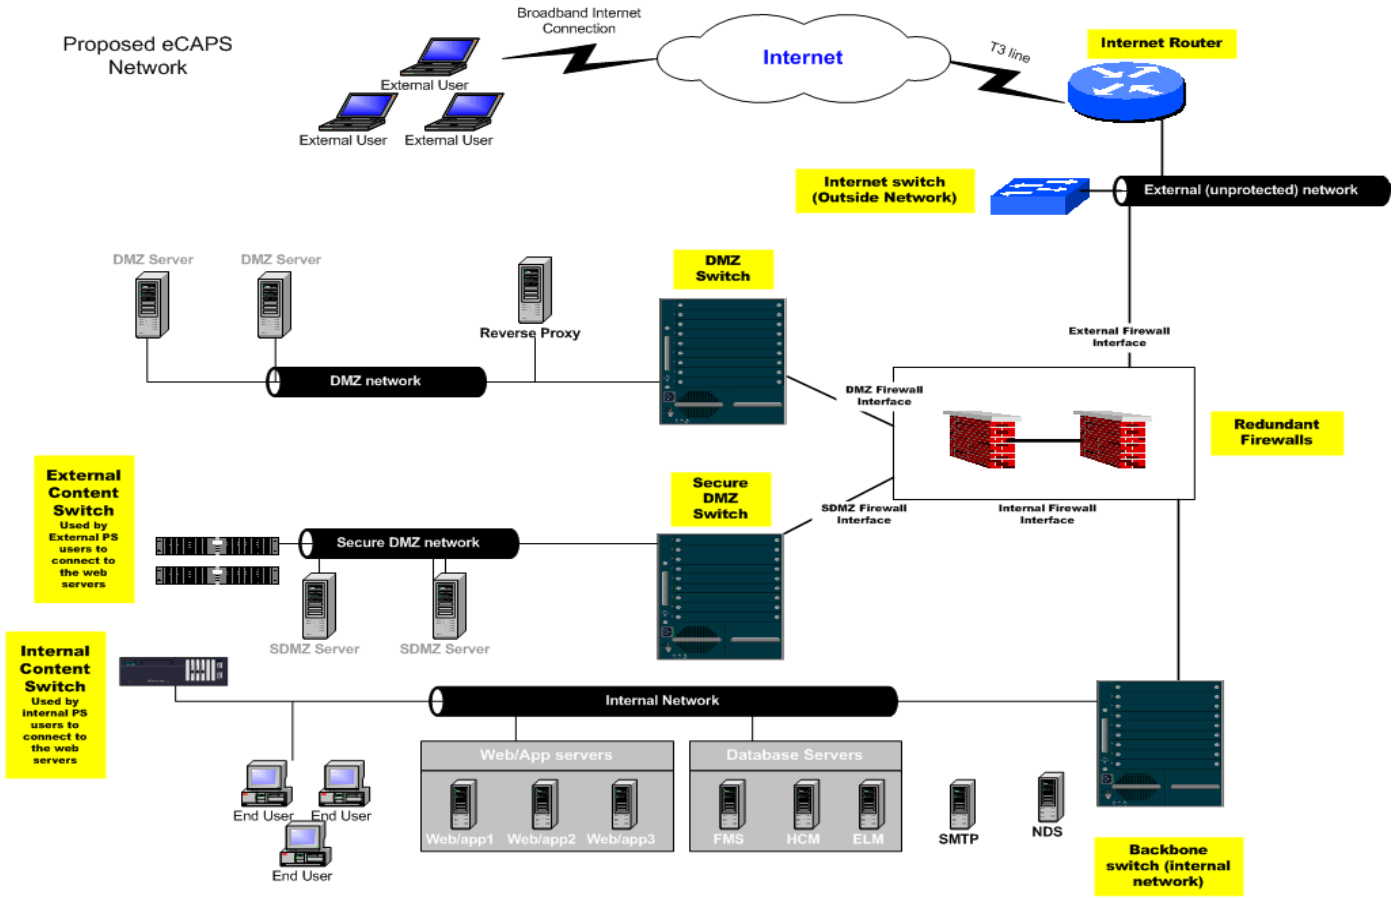

Architecture Design

Load Testing Objectives

- Integrated Workload Testing – take all defined user transactions and mix them into a single workload that is to be representative of a) anticipated production volume, b) peak utilization and c) point of failure.
- Long-term Stability Testing – run an 11 hour stability test to simulate a single workday. This will help identify potential memory leaks and the overall stability or decay potential of the application.
- Batch Process Testing – use scripts to generate volume transactions to be used to simulate nightly batch processing.

Architecture Design

What is Borland SilkPerformer?

- Server-Component Testing tool
- Utilizes Virtual Users (VU's) to simulate utilization
- Script-based testing to replicate online processes
- Consists of SilkPerformer Server and SilkPerformer Agents
- Optional: SAM module to provide detailed metrics

Architecture Design

Borland SilkPerformer Overview

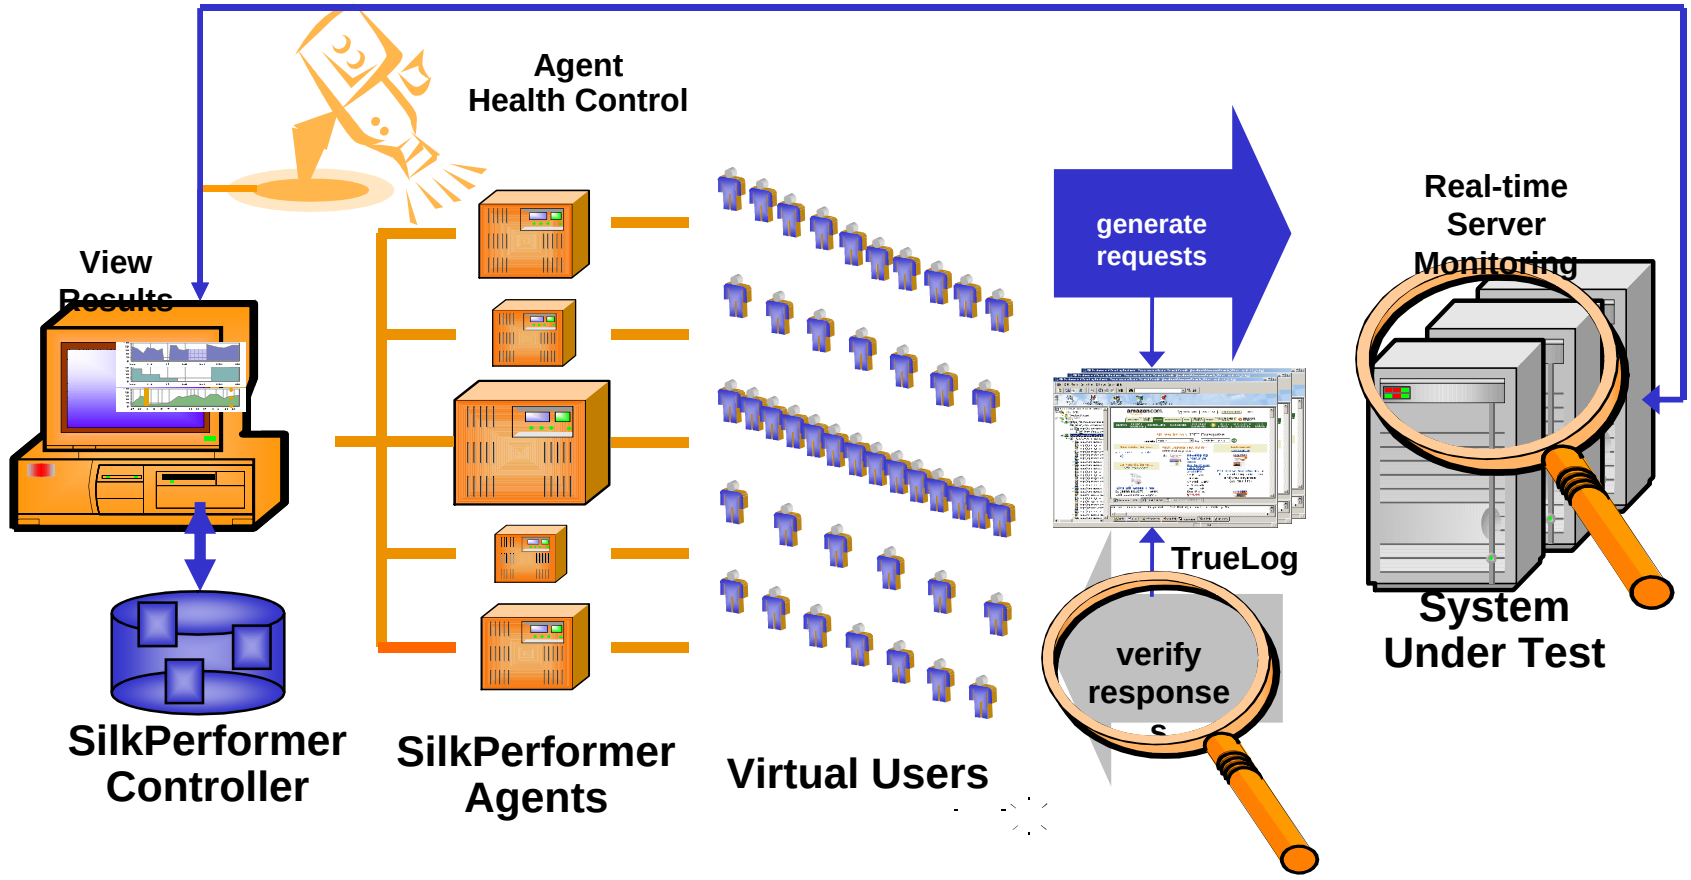

Architecture Design

Prerequisites to Load Testing

- Functional PIA Environment on Production (or production-sized) Servers
- Redundancy
- SilkPerformer Server and Agent Machines
- Load Testing Discovery
- IT Department Support
- Integration Testing
- User Security Profile

Architecture Design

250 Virtual User Load Test

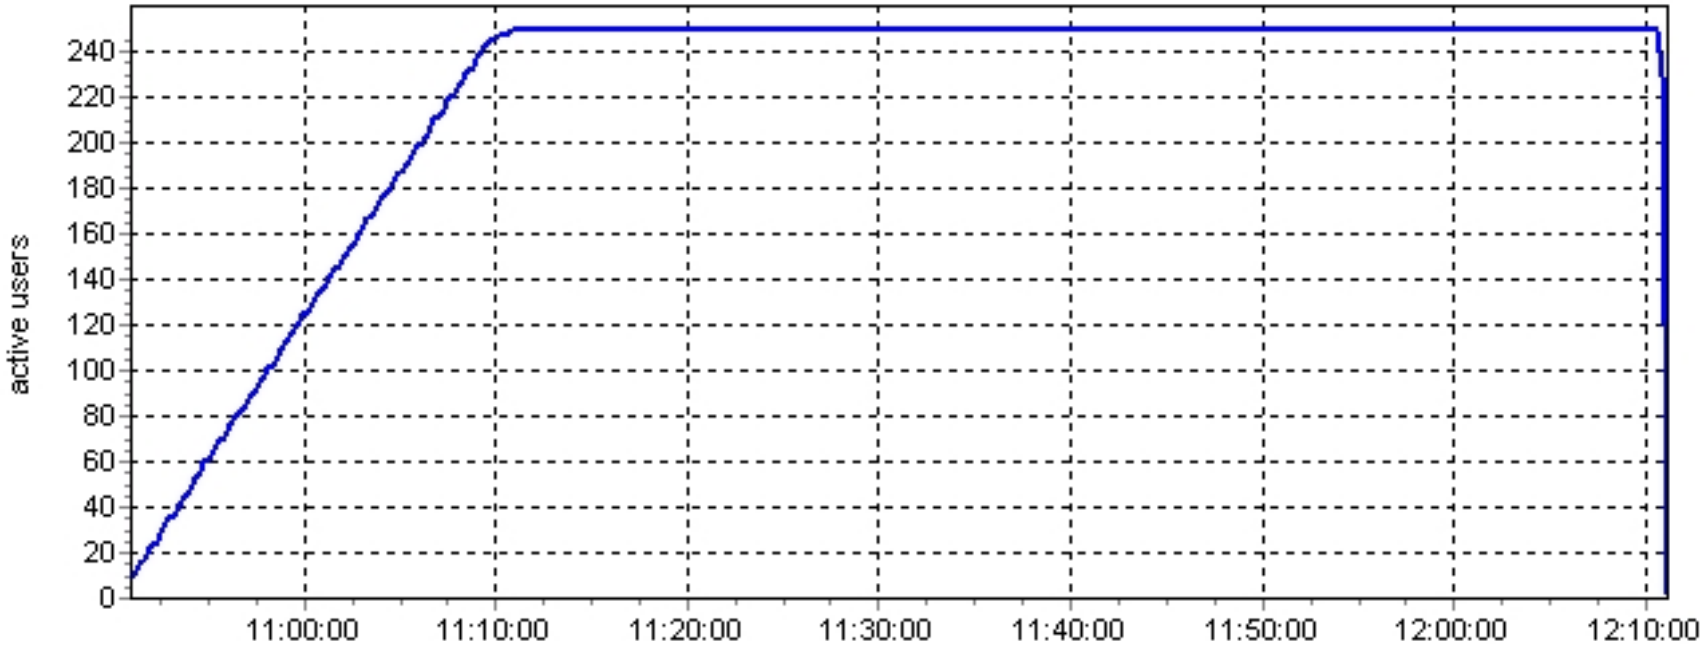

Architecture Design

Load Test Response Times

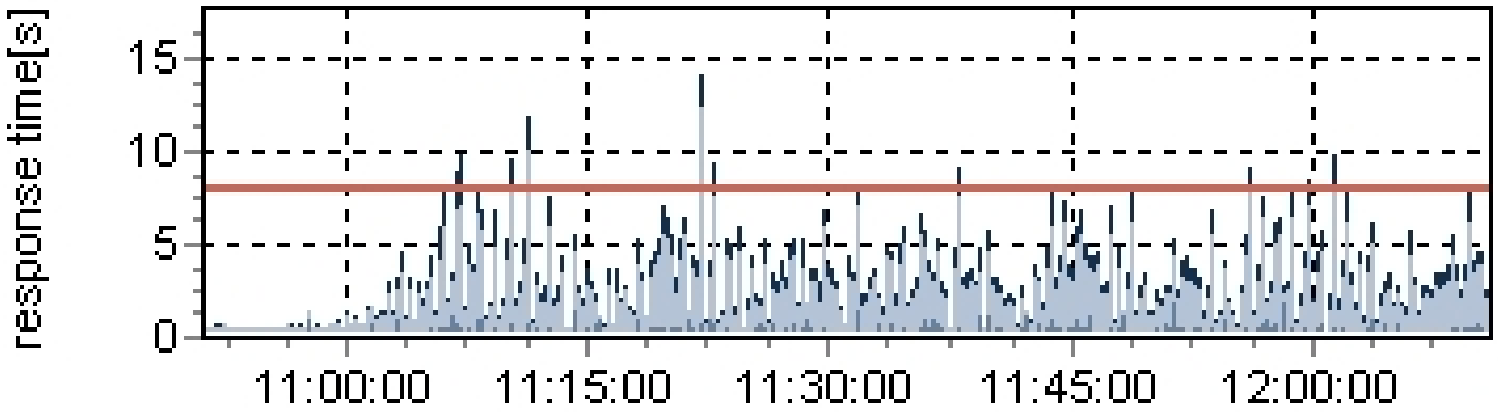

Architecture Design

Performance Tuning Areas

- Database Tuning
- PIA Parameters
- Customization Code
- Batch Processes

Architecture Design

Lessons Learned

- Determine peak utilization requirements early in order to get minimum number of VU's for license.
- Services + spike license is cost-effective option.
- Plan two cycles of Load Testing to incorporate performance tuning in between.
- Disperse SilkPerformer agents across multiple locations to also test different networks.
- Make sure Integration Testing is complete or nearly complete so scripts work properly.
- Define User Security Profile in advance (along with setup click paths) to expedite virtual user generation.
- Install redundancy components prior to Load Testing so entire infrastructure is tested.

Architecture Design

Security

- Secure Sockets Layer (SSL)
  - SSL Accelerator in Content Switch
  - Wildcard Certificate
- De-Militarized Zone (DMZ)
  - Multiple – DMZ + SDMZ
  - Impact on Architecture
    - Content Switch Location
    - Physical vs. Logical Web/App Server

Architecture Design

Multi-Phase Strategy

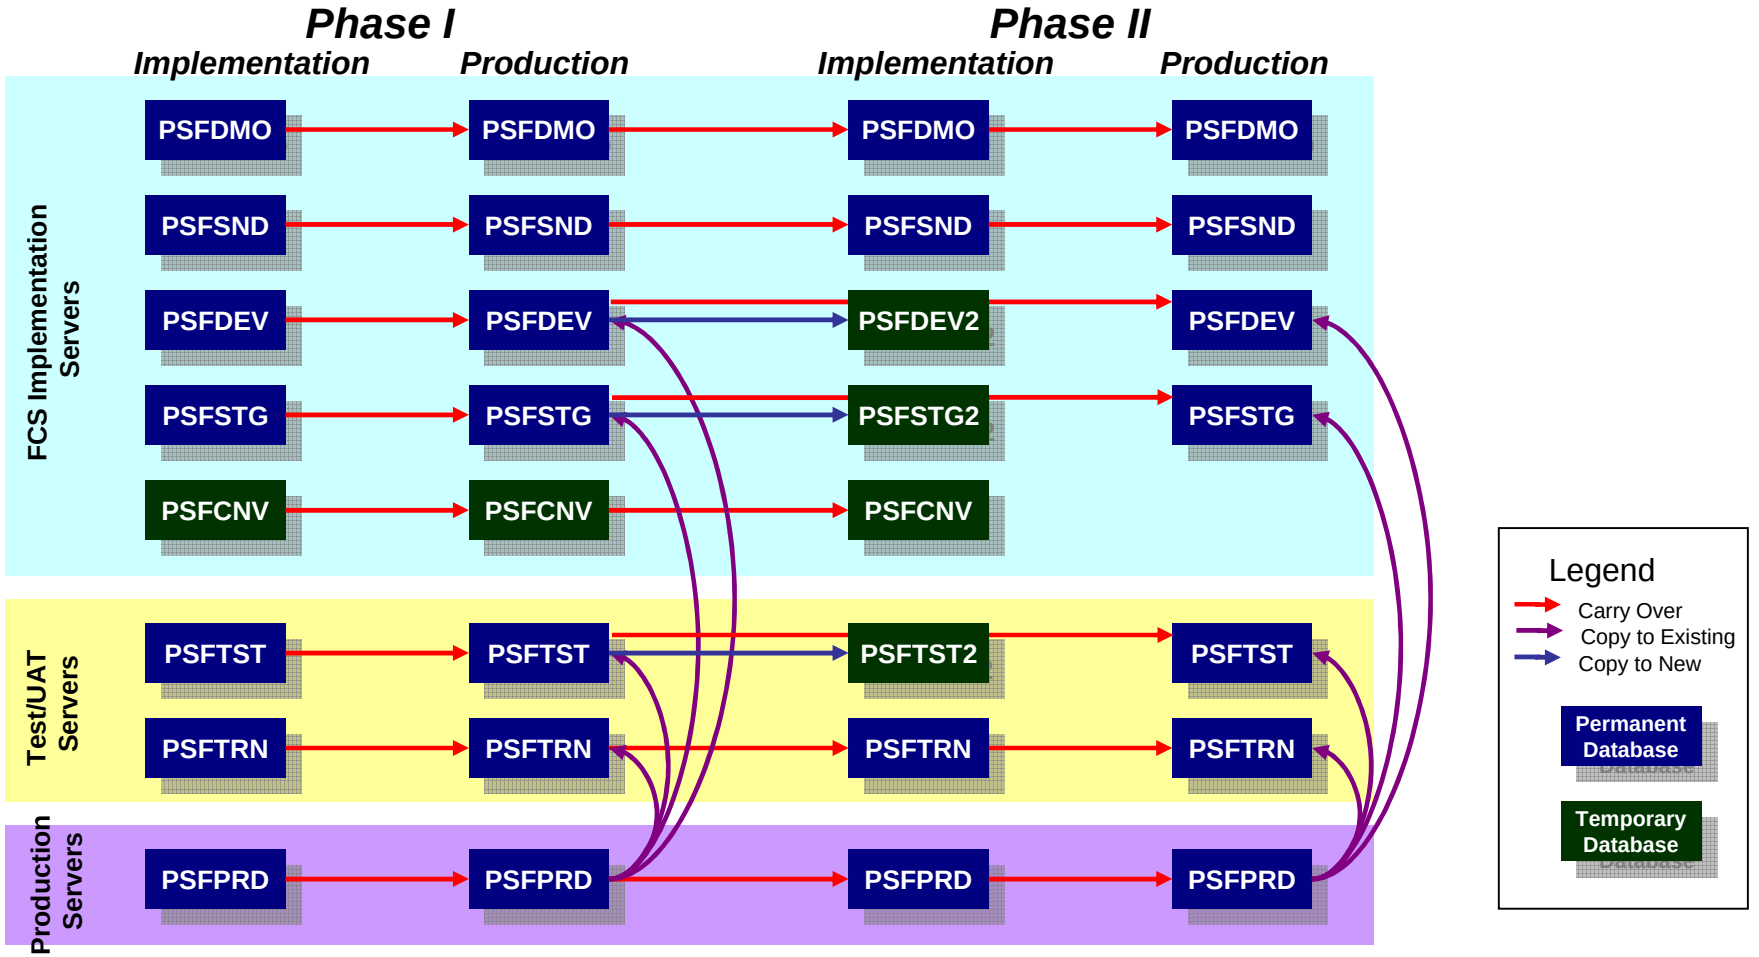

Architecture Design

Lessons Learned

- Redundancy incorporating DR Site requires load balancing/failover across ALL tiers
- Manage eCAPS architecture against non-eCAPS initiatives
- Separate networks can create issues (access, integration constraints, infrastructure management)
- Employee Self-Service inherits back-office access
- LDAP Integration Increases HCM Security Risk
- Load Testing must account for future phases
- Plan Change Control to cover both Production and Phase II activities