Why do systems fail?

Review studies from 1993-98

- Large commercial enterprises

Look to the future

Lisa Spainhower
m eServer Technology

Comparative Study 1993

GENRE	EXAMPLE	DATA SOURCES
Campus-wide LAN	Heterogeneous	Industry Surveys
Mainframe	IBM S/390 9021	IBM Customer Logs Customer Report
Unix SMP	HP 9000	Vendor Claims/IB specs Technical Report Industry Research
HA Mainframe	IBM S/390 XRF	Customer Data
HA Unix	HP SwitchOver/UX	Above Unix Sources Vendor Claims
Fault Tolerant	Tandem Himalaya	Customer Logs Technical Report
????	IBM S/390 Parallel Sysplex	IBM Specs & Models

Comparative Study 1993

GENRE	ANNUAL DOWNTIME (unplanned)	% BY CAUSE HW	SW	other
Campus-wide LAN	453.6 HRS	na	na	na
Mainframe	18.0	4.4	68.3	27.2
Unix SMP	76.1-136.9	34.3	40.7	25.0
HA Mainframe	5.8	0.9	63.7	35.4
HA Unix	21.5	14.9	70.2	14.9
Fault Tolerant	8.9	18.8	74.1	7.1
FT Mainframe	8.4 MIN	14.3	57.1	28.6

1995 Downtime in a poorly-managed S/390 LPAR

	Impact events	Events	
ATTRIBUTION	# OUTAGES	# OUTAGES	IMPACT RATIO
Control Center	70	24	2.9
Environment	18	5	3.6
Hardware	10	1	10
Software	118*	52	2.3
Total	216	82	2.6

ATTRIBUTION	OUTAGE (min)	OUTAGE (min)	IMPACT RATIO
Control Center	5202	1949	2.7
Environment	1275	454	2.8
Hardware	875	88	10
Software	6209	3062	2.0
Total	13561	5553	2.4

*TM-56%, Apps-16%, DBA-14%, OS- 6%, other-

1995 Downtime in a poorly-managed S/390 LPAR

Total Outage per log: 226 hours
Per one outage/event: 93 hours

✓# 1 Contributor is software: product & process

1818 process
453 product
791 uncertain
3062 total (51 hours)

Assume all CC outages are process (1949 min.; 32 hours)

✓68%-82% of all unscheduled outages are process
(63-76 hours)

✓Technology - HW/SW - 10-24%
(9-22 hours)

Aggregated UNIX server data

Downtime Cause %	UNIX Standalone	UNIX Cluster
Hardware	42	46
Software	34	36
Other	24	18

Data from a very large well-managed Unix customer

% Unplanned downtime	
HW	43.8
OS	7.8
App	7.3
Com link	18.8
DB	2.0
Environment	0.7
Supplier	0
Op tools	1.8
Process	6.2
Org/structure	0.8
Human error	5.0
Other	5.8

% Planned downtime	
HW maint	26.5
OS install	1.6
App Release	34
Com link	2.5
DB admin, BU	31.5
Dis Rec Test	0
Pwr Test	0.2
Other	3.7

1998 Unix Investigation

GG claims Unix HA/FT clusters limited to 99.9% availability until 2000.

HP has 99.95% guarantee; announced plan for 5 9s in 2000.

Objective

Determine achievable availability for Unix for 98-00

Major Limitation

Very limited, inconsistent data available

1998 Unix Investigation

	Customer	DH Brown
# clusters/nodes	11/63	15/38
unsched. avail.	99.6%	99.99%
failover/yr/node	2.4	1.4
downtime/failover	>1 hr	15 min
node failover	-	4-20 min

- With improved HW: 8.6-76 min/yr/2 nodes
- Software retry will also improve
- Measured as HP specifies, 4 9s is feasible with 5 9s for OPS and Web server

Lessons from the 90s

- Management discipline is critical to HA
- Fault tolerant servers make a difference
- Clusters are difficult to implement

Making the Front Page

Source: Gartner Group

eBay
12 June 1999 outage: 22 hours
Operating System Failure
Cost: $3 million to $5 million revenue hit
26% decline in stock price

AT&T
13 April 1998 outage: Six to 26 hours
Software Upgrade
Cost: $40 million in rebates
Forced to file SLAs with the FCC (frame relay)

America Online
6 August 1996 outage: 24 hours
Maintenance/Human Error
Cost: $3 million in rebates
Investment: ???

Causes of Unplanned Application Downtime

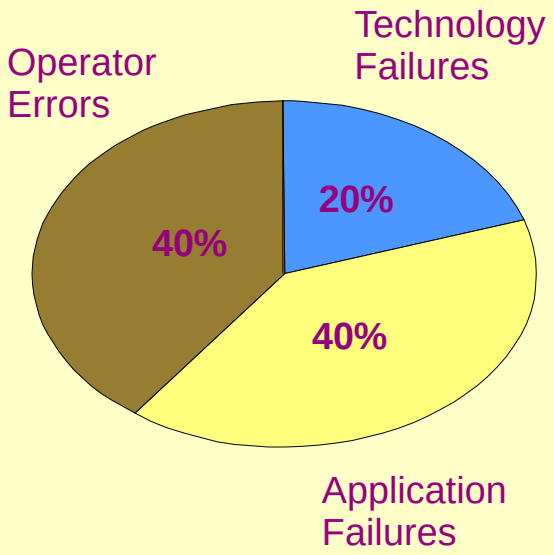

Charles Schwab & Co.
24 February 1999 through 21 April 1999: Four outages of at least four hours
Upgrades/Operator Errors
Cost: ???; Announced that it had made $70 million in new infrastructure investment. $s

Dev. Bank of Singapore
1 July 1999 to August 1999:
Processing Errors
Incorrect debiting of POS due to a system overload
Cost: Embarrassment/loss of integrity; interest charges

E*Trade
3 February 1999 through 3 March 1999: Four outages of at least five hours
System Upgrades
Cost: ????
22 percent stock price hit on 5 February 1999

Managing Exploding e-business Infrastructure

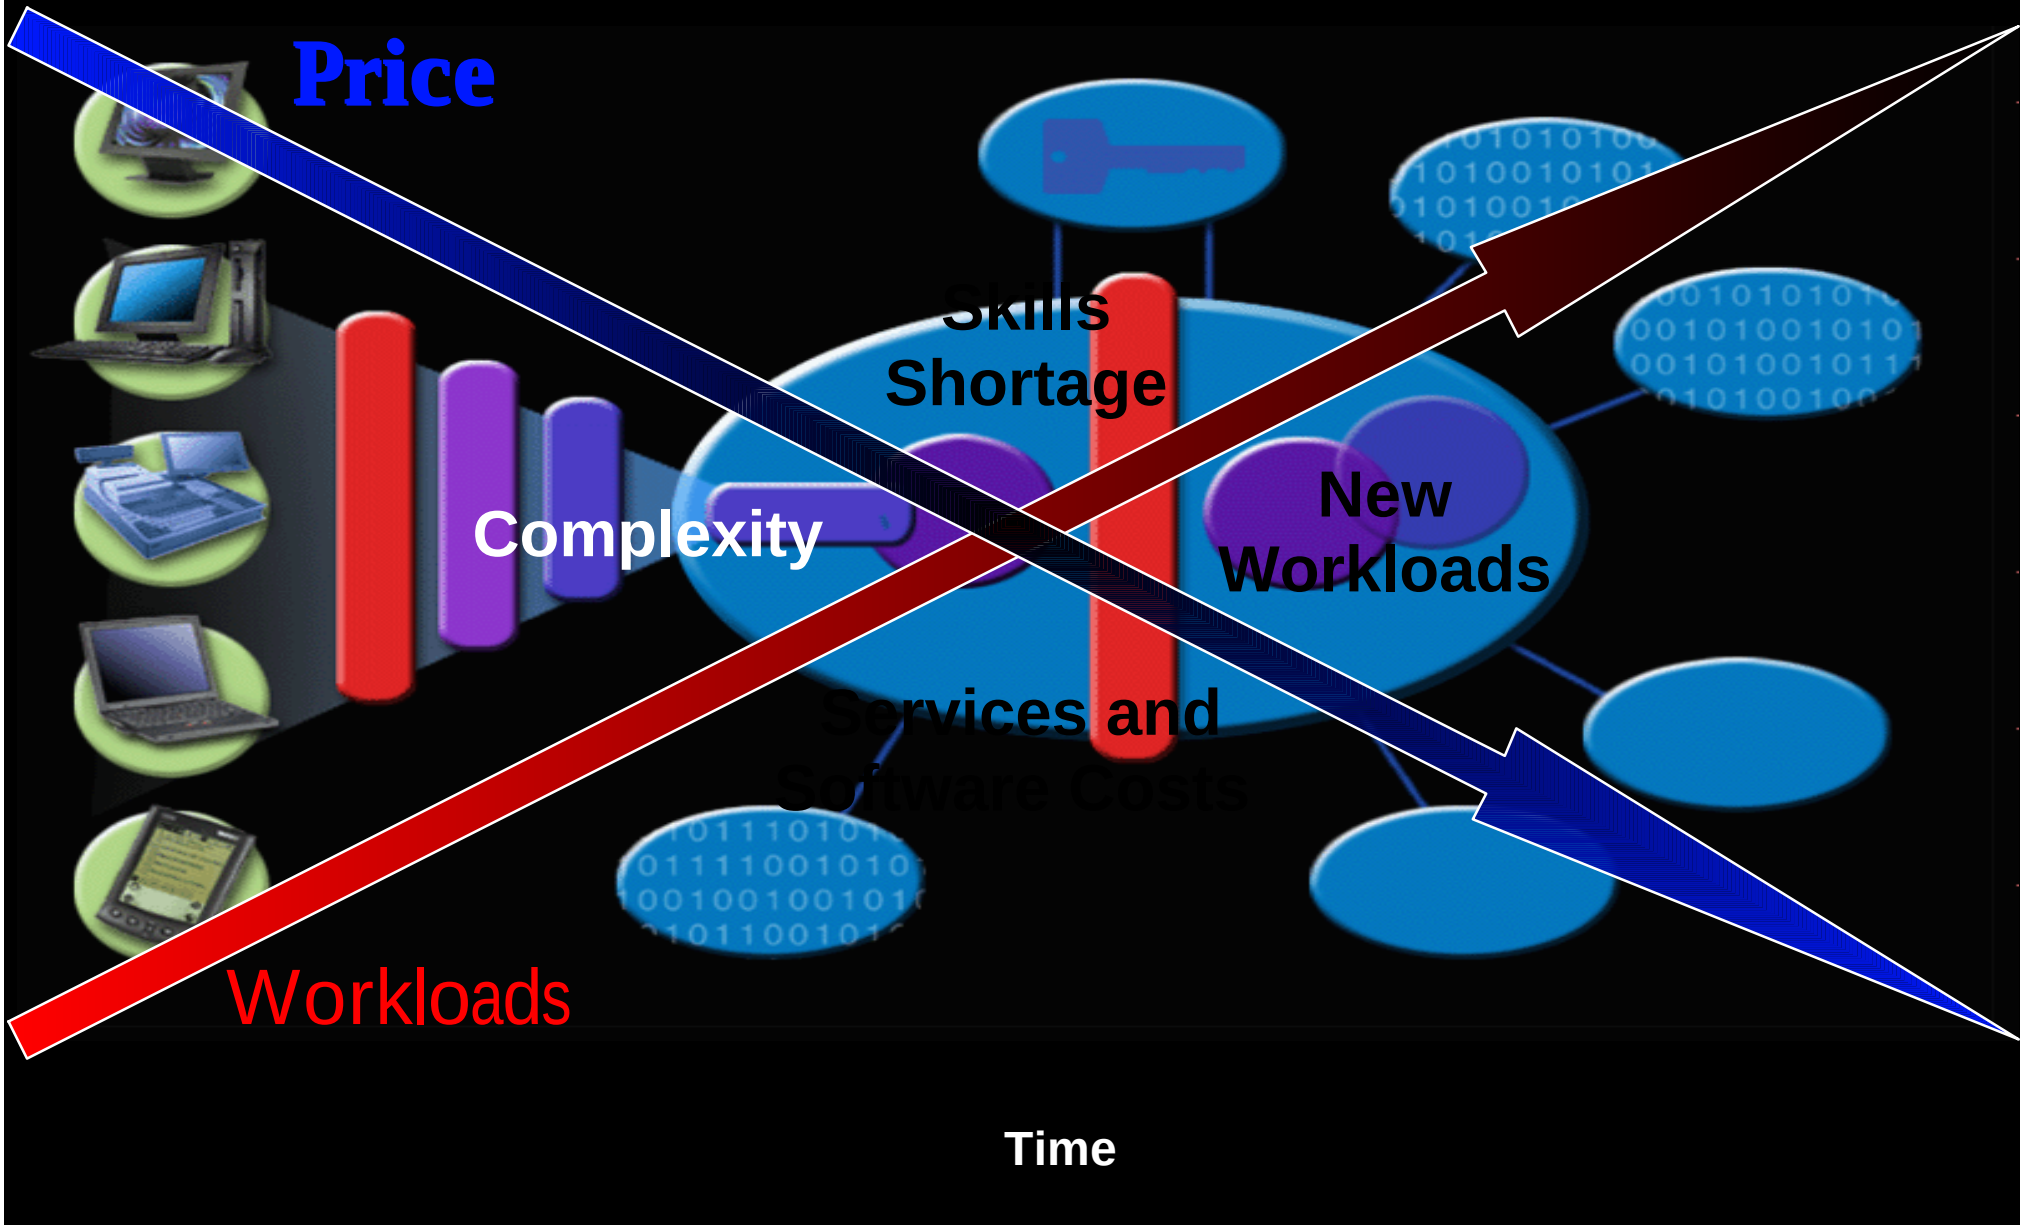

Challenges for the 00s

- Increased importance of firmware
- Circuit failure mechanisms
- State encapsulation
- On-the-fly change
- Dynamic resource allocation
- Configuration validation