Workshop on Dependability of e-Business Systems
Tuesday June 27, 2000
Summary and Notes on Workshop

Nick Bowen
June 28, 2000

Workshop on Dependability of e-Business Systems
Tuesday June 27, 2000

- Goal
  - Create a lively discussion among
    - Companies deploying major eBusiness systems (>$1B/year)
    - Companies building system technology
    - Academics doing Research in this area
  - Attracting a number of people from outside FTCS community
  - Leading to tangible output around a common[1] view on key future issues

[1]common ^= agreement

Workshop on Dependability of e-Business Systems
Tuesday June 27, 2000

- Program Committee
  - Nick Bowen (IBM Research)
  - Wendy Bartlett (Compaq Computers)
  - Linda C. Ernst (Intel)
  - Michael Treese (Sun)
  - Lisa Spainhower (IBM)
  - Brendan Murphy (Microsoft Research)
  - Bob Horst (3ware)

Workshop on Dependability of e-Business Systems
Tuesday June 27, 2000

- Guidance given to panelists
  - Environment
    1. What is your business environment and what are the key workloads?
    2. How do you define "high availability or fault tolerance" and what is the cost of outages both in monetary terms and in terms of brand image?
    3. *What failures do you plan for, do these include off-site disaster recovery?*
    4. What is your technical architecture?
  - Experiences
    5. *What are your experiences, i.e. what have you learned? For example, how do system failures affect end user availability or how to disk failures or maintenance items like database re-orgs affect your availability.*
    6. How do you see your environment, both business requirements and technical evolution, changing over the next five years?
  - The Future
    7. What message do you have for industry on what product features they should be building?
    8. What are the really hard problems where you think the Research community should be making new inventions?
    9. What is your view of the reliability of the internet and what impact does this have on you?
    10. What features/inventions need to be added to the base internet infrastructure?

Final Agenda

8:30 - 9:30	"Internet Performance/Availability from an end user perspective" Eric Siegel, Keynote Systems
9:30 -10:00	Break
10:00 -12:00	Panel I: Dependability for Business to Business Systems, Chair: Brendan Murphy
12:00 - 1:30	Lunch
1:30 - 3:00	Panel II: Dependability for Consumer Based Internet Systems, Chair: Michael Treese
3:00 - 3:30	Break
3:30 - 4:30	"Losing Sleep After the Hardware Issues are Solved," Brian Brandt, Strategic Consulting, Inc.
4:30 - 5:30	Forum: Key Issues for Industry, Research and System Development, Chair: Nick Bowen

Attendance

Full Room	"Internet Performance/Availability from an end user perspective" Eric Siegel, Keynote Systems
60 - 70 people	Panel I: Dependability for Business to Business Systems, Chair: Brendan Murphy
32 people	Panel II: Dependability for Consumer Based Internet Systems, Chair: Michael Treese
60 people	"Losing Sleep After the Hardware Issues are Solved," Brian Brandt, Strategic Consulting, Inc.
55 people	Forum: Key Issues for Industry, Research and System Development, Chair: Nick Bowen

Invited Talk (8:30 - 9:30)

- Internet Performance/Availability from an End User Perspective, Eric Siegel, Keynote Systems
  - This talk provides a dose of reality by showing a view of Internet performance and availability as seen from an end user's perspective. The material is based on extensive measurements of real world web sites. The company captures more than 16 million performance measurements daily using Keynote's global infrastructure of over 220 measurement computers connected to the major Internet backbones in over 90 statistically selected Internet access locations across 45 metropolitan areas worldwide.
- Summary
  - full group attended
  - Large backlog of questions
  - Good tutorial of internet, especially aspects that cause real problems (e.g., peering networks)
  - Good transition to real issues with story about them monitoring 25 brokers from 120 monitoring locations. There failure rates (daily?) Range from 0.2% to 30%.
  - Statistics of failures interesting, not central system failures.
  - Need to measure total user experience.

Panel I

- Panel I: Dependability for Business to Business Systems, Chair: Brendan Murphy, Microsoft
  - Bradley A. Ellison, Intel
  - Tony Lostaglio, IBM
  - John Steensen, Intira
  - Tom Buiocchi, Peakstone
- Technical details
  - 60-70 people attended
  - 74 minutes of panel talks (we started 10 minutes late)
  - 30 minutes of questions

Panel I - Key Points from panelists

- John Steensen, Intira (Infrastructure for ebusiness outsourcing; 1998 -)
  - huge skill set required in real world: Servers, DB, network, security, operations
  - Humans must still watch the system
  - Most real systems have over 1000 parts and cannot be modelled
- Tom Buiocchi, Peakstone (1 year old company (no product for 45 more days))
  - Guarenteed Service Levels
  - Differentiated services
  - Resilent to spikes and DoS attacks
  - Highly automated
  - Precision Provisioning
- Bradley A. Ellison, Intel (Plan is for 95% of 2000 revenues over the net)
  - Over 500 servers: 33% production, 66% test & development
  - 99.8% availability - 86 minutes
  - New software installations are major cause of failures
  - Need distributed systems - how to distribute Dbs?
  - Things to fix:
    - Application like ease of use (plug and play)
    - Performance
    - Increase manageability
- Tony Lostaglio, IBM ($20B of ecommerce in 2000 expected)
  - 4000 new "web shopping experiences" will deployed on te web in 2000 - Pretesting?
  - End to end is a massive collection of old + new systems
    - How do you decide how to invest (for availability) in each one?
  - End to end availability is 92-93% with individual systems at 99.5 - 99.9%

Panel II

- Panel II: Dependability for Consumer Based Internet Systems, Chair: Michael Treese, Sun Microsystems
  - Carl Hutzler, AOL
  - Brian Koster, Dell
  - Richard Gornitsky, PaineWebber
- Technical details
  - Started 10 minutes late because of post-lunch low attendance
  - About 32 people attended
  - Panelists talked for 65 minutes, 15 minutes of questions

Panel II - Key Points from panelists

- Dell (Kevin Koster)
  - $0 in 1995, $1M / day on net in 1997, $45M/day today!
  - Run web site "all on dell" (many refurbished systems)
  - Things they do:
    - Geographic redundancy
    - Multiple ISPs
    - Rigorous development methodology
    - Application clustering
    - Limit access to production system
  - Detection
    - Infrastructure monitoring
    - Application monitoring
    - Application instrumentation
  - Biggest challenges
    - Managing change
      - Understanding the impact of change
      - Testing, when there is no time to do so!
      - Test bed
    - Monitoring the customer experience
    - Managing capacity
      - Predicting capacity effects of new business capability
      - Mapping customer behavior to infrastructure capacity
  - Monitoring infrastructure is very costly, you must make business tradeoffs
    - They gather 1 TB of monitoring data per quarter

Panel II - Key Points from panelists

- AOL (Carl Hutzler)
  - "email" is their #2 mission critical application (behind "login")
  - Dealing with 2x growth of system every 11-13 months
  - Tandem NSK systems core core email (50 systems, record is 134 million email postings in 24 hours, with 72 in one hour.)
  - 35 TB+ of email content
  - Server and data redundancy
  - Strong QA & production methodlogy
  - Planned downtime decreased 10x in the last year
  - Unplanned outages are still low (99.996%, up, excluding planned)
  - Future needs
    - Availability: ONLINE maintainence - OS, DB, etc.
    - Reliability: testing of new SW-HW needs to improve
    - Disaster recovery: Replication of data
- PaineWebber (Richard Gornitsky)
  - HA Drivers
    - Components failures must be masked from end users
    - Ensure responsivenss off and on trading hours
  - 12 components, each 98.5 - 99.9 -> ~90% availability
  - HA study they did
    - HW is reliable, humans are not
    - Applications need to be implemented with change & problem management in mind
    - Apps need to handle exceptions and perform fail over / recovery
    - Time to market is incompatible with HA
    - Instrumention is critical
    - Cost of HA may not be worth the investment
    - Industry must invest in software engineering techniques

Invited Talk

- Losing Sleep After the Hardware Issues are Solved, Brian Brandt, Strategic Consulting, Inc.
  - This talk highlights several of the keys issues that are facing businesses as they enter the 21rst century. Particular focus is paid to those issues around providing high availability in an increasingly complex world where customer demands are constantly increasing.
- Summary
  - Had about 60 people
  - Back end connectivity is a big problem
    - diversity of backends
    - multiple network hops
    - Insanely code and time sensitive
    - Suggested invention needed here
  - Storage scalability
    - Suggested need to invest in massively scalable storage unit (in small units)
  - Win2K monitoring and management a problem
    - Suggested ways to fix this, but still a problem

Interactive Session

- Key Issues for Industry, Research and System Development , Chair: Nick Bowen, IBM Research
  - During this session participants have the opportunity to express their opinions on the critical issues that face industry, the research community and people deploying these systems. We expect a fruitful debate among this community on where the real issues lie.

Interactive Forum

Workshop on Dependability of e-Business Systems

Goal: Get a perspective from people in the e-Business trenches so that the Research community and Product developers understand the real life issues.

4:30 - 5:30 Session:

- Create 1 slide with your opinion on the top issue (plus name)
- Take one clarifying question (if one is out there)
- Free drink (on the boat) for first 14 people!

Interctive Forum - Summary

- 15 people with direct contributions
- 22 unique ideas
- lot's of discussion within the room (often had three mic's discussing)
- lasted 45 minutes
- About 55+ people attended the whole time (till 5:15).
- Mostly open questions and observations

Interactive Forum - Summary

- Testing
  - Systems are so big that you cannot have a real test environment
  - Time to market of new apps <<< traditional test time and tools
- Predicting Capacity
  - Predicting the load impact on the overall system of a new application is practically impossible.
  - User behavior can change everything (browse v. Buy)
- 99.9^N Effect
  - Systems are complex and consist of many (sequentially accessed) components (>10), often indepedently owned which each have ~99.5-99.9 availability.
- E-ngineering - precious little engineering discipline, ok now, but not when we get to life critical applications.

Interactive Forum - Summary

- Management Issues
  - Systems Management
    - Are new management paradigms needed to handle the huge growth in the number of components at the same time there are fewer & fewer skilled people?
  - Design for Manageability
    - Are application developers going to supply the needed instrumentation?
  - Change Management
    - Conflict between need for change and problems introduced by change.
  - System / Application fault management
  - Dependability Measurement
  - Scalable Management Solutions
    - instrumentation of infrastructure ("measure, measure, measure")
    - collaboration across disperate distributed systems
    - abstraction / correlation to a higer level , e.g., FRU precision

Interactive Forum - Summary

- Security Issues
  - We must figure out how to stop criminals, competitors, and idiots from using your network against you.
  - Why the lack of discussion by panelists on intrusion attacks - Not a problem or too sensitive to discuss in open?
    - Started long discussion, e.g., Dell panelists - clearly a big issue (at least one attack a day), not an availability problem, yes, though, hesitent to expose this in public.
  - Security Issue: How do you prevent hackers from stealing mastercard account numbers?
  - 10 years of security works seems to have been ignored - firewalls and ad hoc tools, industry has decided to go ad hoc way.

Interactive Forum - Summary

- Disconnect between Academia and the real world
  - Problem today: Academic(too theoretical), Government (too slow), Business (too busy)
    - Need an "E-Business Foundation" (aka. Software Eng. Institute) with balance of Resaerch, Education and Development
  - Gap between Research and Practice is (appears to be) growing
    - Is there a way to guide and incorporate results into real problems?
    - New centre?
  - Problems for Universities vs. Internet Technology
    - Industry moving too fast
    - Ad hoc solutions to Ad Hoc problems
    - Poor literature pool
  - There seems to be huge problems with e-Business systems but it is not clear what the Research issues are

Summary

- Getting these industry panelists was **very** hard.
- Lots of very positive comments
- Large core of people attended the whole day
- Lots of good insights came out of the day