

Reboot-based High Availability


Turning evil reboots into reliable friends

George Candea Armando Fox
Stanford University

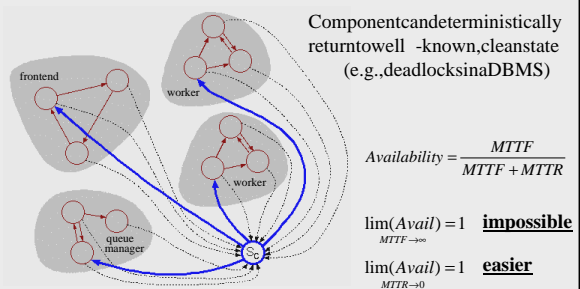
Motivation

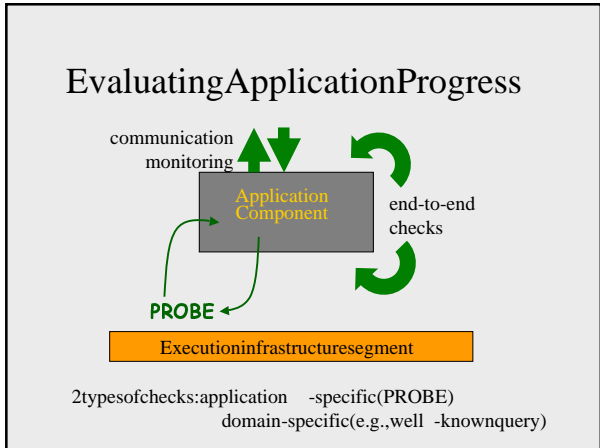
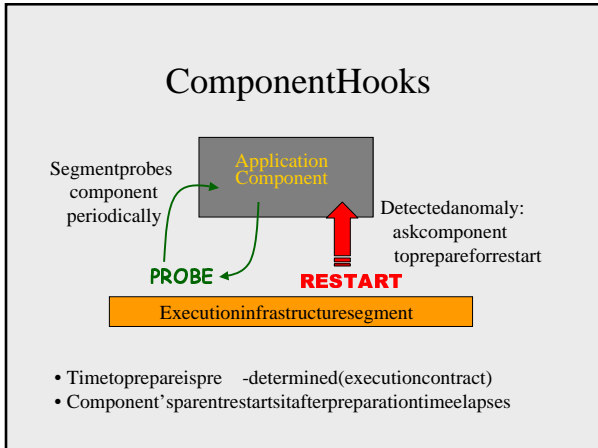
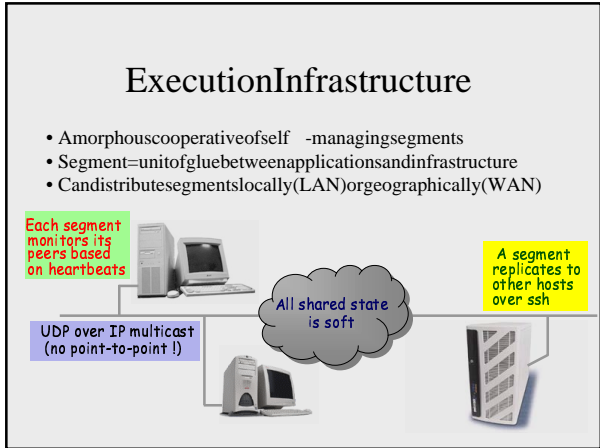

- Software infrastructures = aggregation of industrial strength software components, yet...
 - 40% of total unplanned downtime due to application failure [Gartner]
 - Business losses: average \$6.45 million/hour of downtime for brokerage industry [Dataquest]
- Good news: $\geq 90\%$ of bugs in production -quality software are transient [Adams, Gray] → **REBOOT!**
- System management costs \gg installation costs

Modeling Software Systems



Modeling Partial Reboots





Applications as Distributed Systems

- Functional distribution
Distribute components, even if logically colocated
- Loose coupling
Glue components together with announce/listen protocols
- Minimal inter-component assumptions
When components make implicit assumptions, they are overflowing their boundaries
- Weaker guarantees, stronger best effort
E.g., IP is extremely robust, in spite of not guaranteeing much

Restartable Software

- End, don't grant
Simplifies recovery and restart because failed system returns to a clean state after the lease times out
- Persistent state → soft and/or degradable state
Trade consistency for availability
- Orthogonal mechanisms with minimal state sharing
Maximizes effectiveness of partial reboots

Evaluation Methodology

- Compare to existing high availability mechanisms using fault injection
- Evaluate execution infrastructure at different levels of application modification
 1. No changes
 2. Rudimentary PROBE and RESTART
 3. PROBE/RESTART + restartability guidelines
- Deploy in 24x7 environments

Related Work

- Use periodic reboots to prevent failures caused by software aging*
Y. Huang, C. Kintala, N. Kolettis, N. D. Fulton, Software rejuvenation: analysis, module and applications. Proc. FTCS 1995, pp. 381 - 390.
- Distributed execution platform with segments that migrate in response to failure*
J.F. Shoch, J.A. Hupp, The "Worm" Programs - Early Experience with a Distributed Computation, CACM 25(3): 172 - 180 (1982).
- Generic, transparent application recovery is impossible in most cases*
D.E. Lowell, S. Chandra, P.M. Chen, Exploring Failure Transparency and the Limits of Generic Recovery, Proc. OSDI 2000.
- How to write software components that are easy to reuse and compose*
D. Garian, R. Allen, J. Ockerbloom, Architectural Mismatch or 'Why it' shard to build systems out of existing parts, Proc. ICSE 1995, pp. 179 - 185.

References

- [Adams]
E.Adams,Optimizingpreventativeserviceofsoftwareproducts, IBMJournal
ofResearchandDevelopment,28(1):2 -14(1984).
- [Dataquest]
J.Sheridan,HighAvailability – HowHighCanYouGo?,Dataquest
TechnologyAnalysisPerspective, Gartner Group,September1996.
- [Gartner]
D.Scott,MakingSmartInvestmentstoReduceUnplannedDowntime,
ResearchNote, Gartner Group,March1999.
- [Gray]
J.Gray,WhyDoComputersStopAndWhatCanBeDoneAboutIt?,P roc.
SRDS1986,pp.3 -12.