



Automating Disaster Recovery: the ultimate reliability challenge

SREcon24 Americas

Ricard Bejarano

Lead Site Reliability Engineer, Infrastructure

Something is up with the system



What is up with the system?



What could've been done
differently *a decade ago* that
would've prevented this?



Most incidents are
resolved (prevented)
at the drawing board



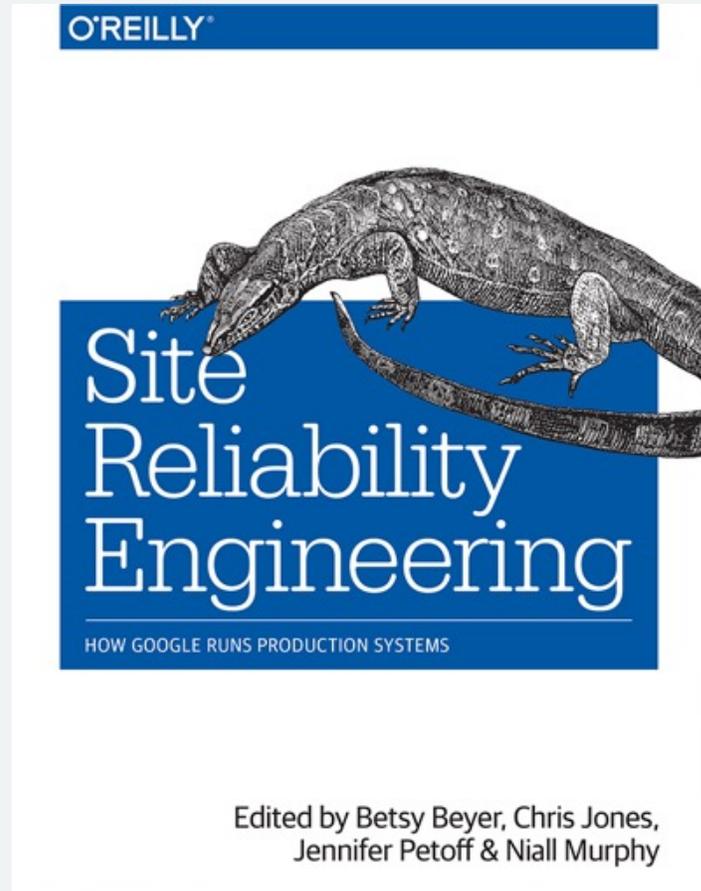
So... what could've been done
differently *a decade ago?*



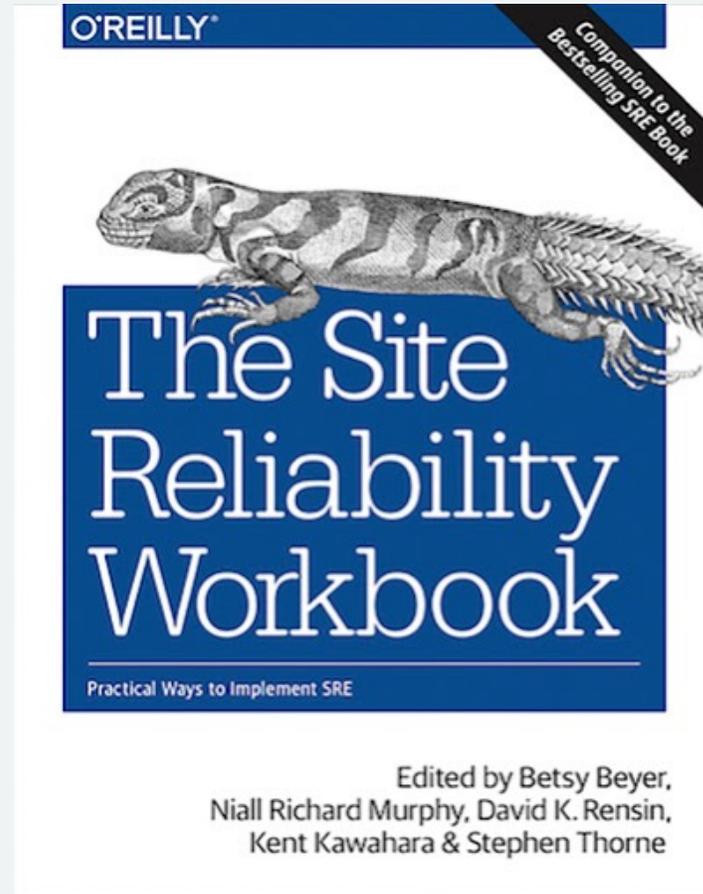
Reliability cannot be an afterthought



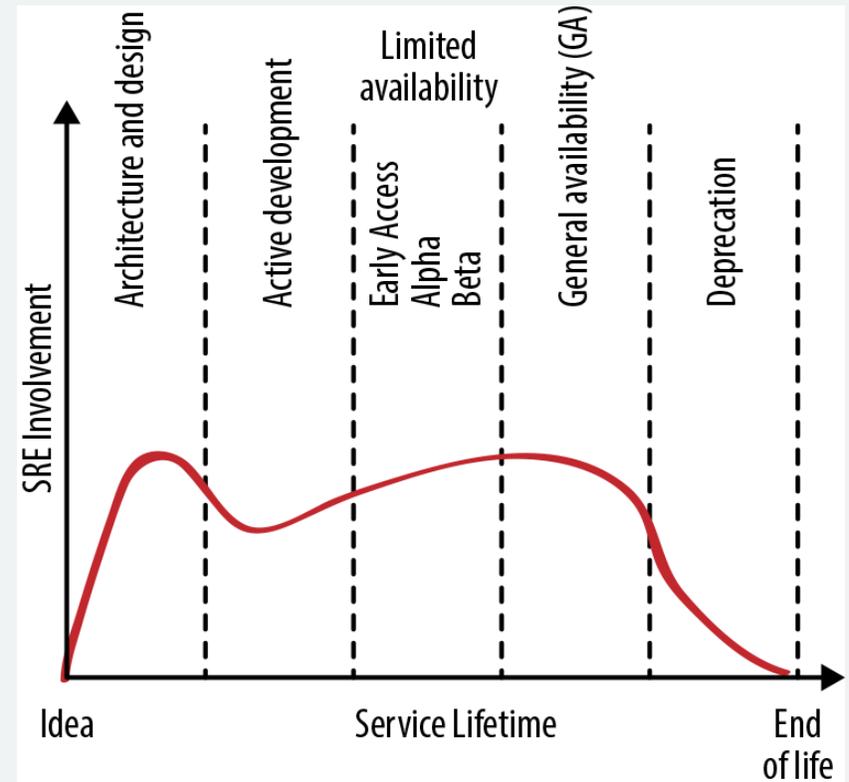
Who has read
this book?



Who has read
this other book?



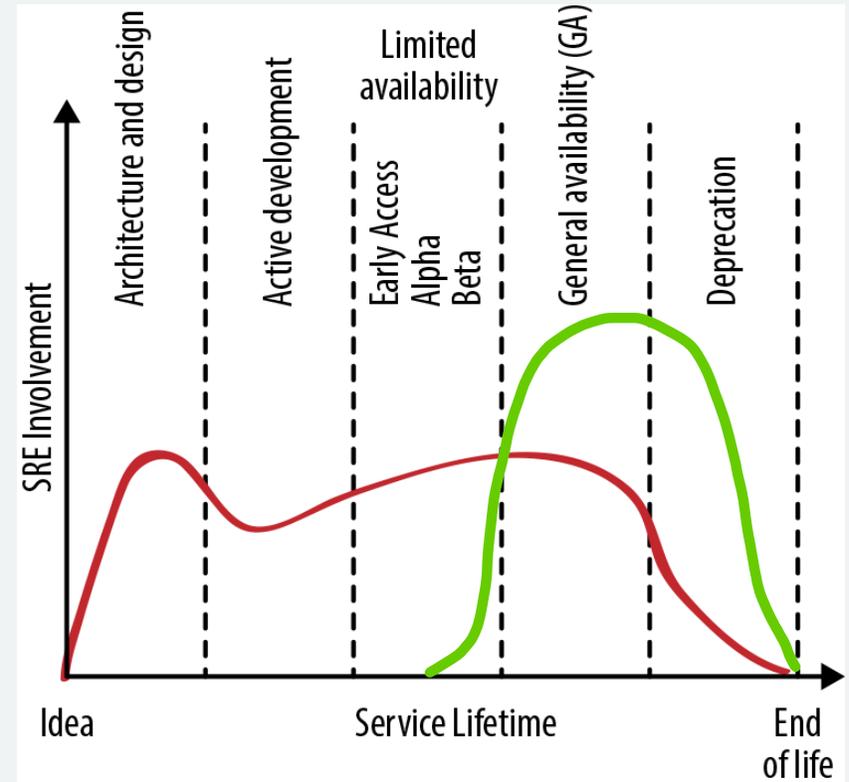
Engagement over lifetime



The Site Reliability Workbook
Fig. 18-1



(actual) Engagement over lifetime



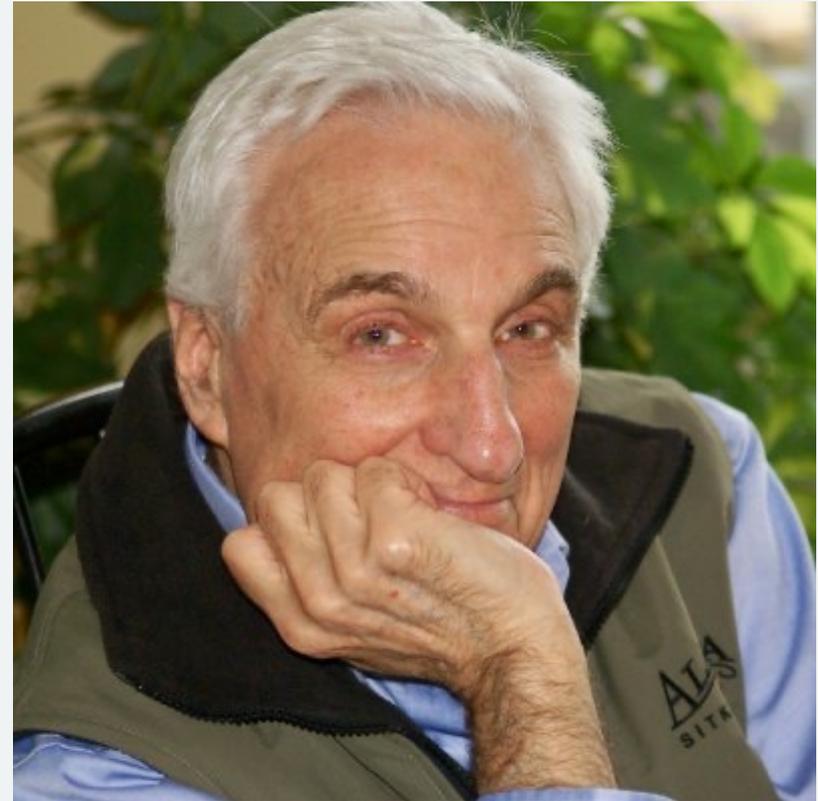
The Site Reliability Workbook
Fig. 18-1 (modified)



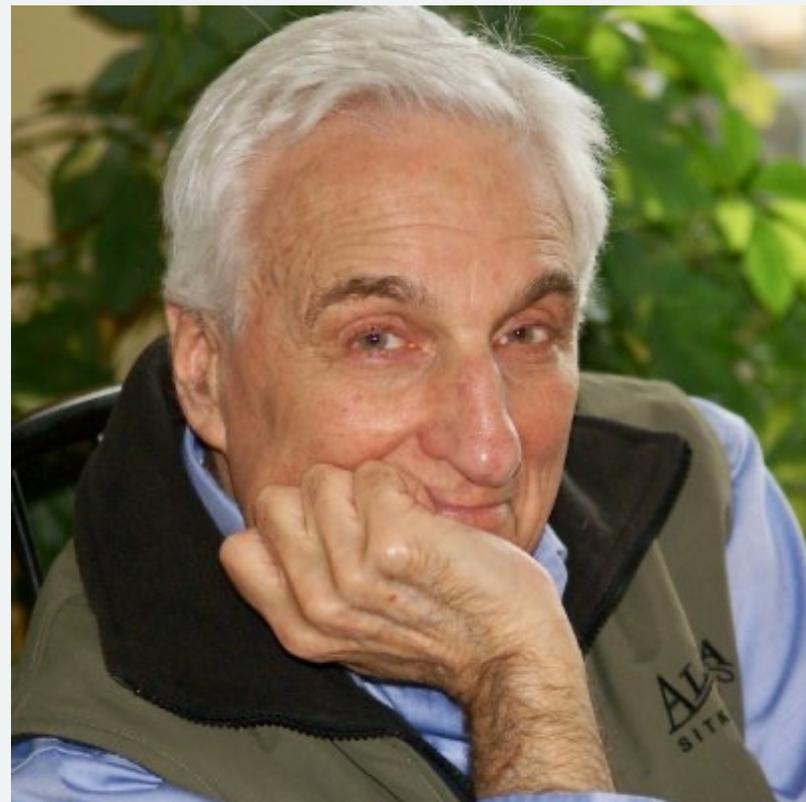
Reliability cannot be an afterthought



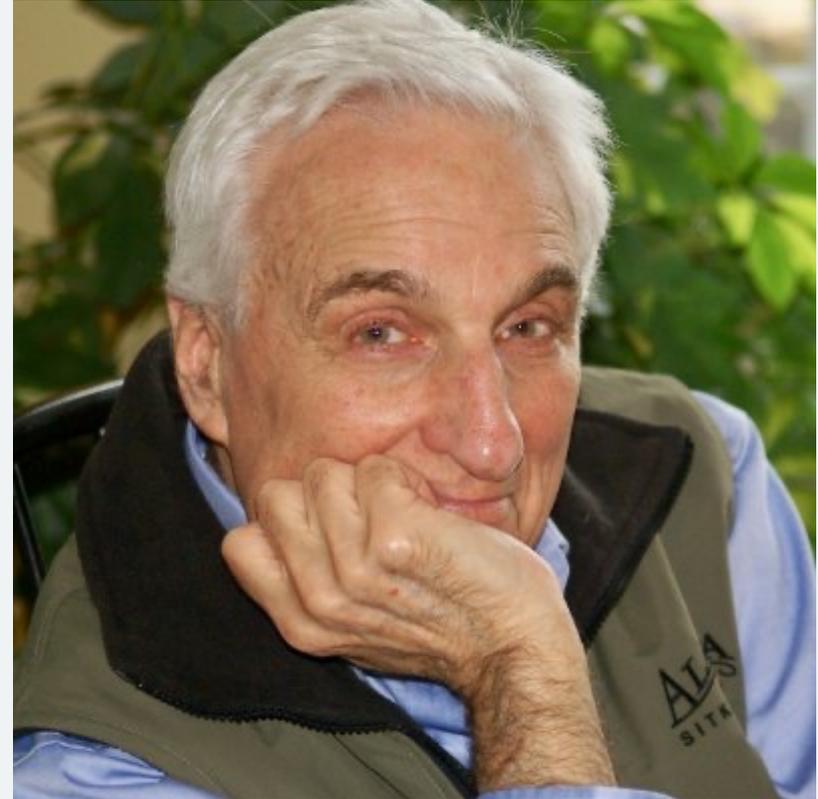
Who's this?



Melvin Conway



Melvin Conway's law



The architecture of the software you're building, will closely resemble that of the human structure you set out to build it



Should SREs influence
team structure decisions?



So... what could've been done
differently *a decade ago?*



Putting SREs in the room **earlier**



Reliability-driven development



What about automating
system recovery?



Priority number 1 is still *recovery*



Generic mitigations

oreilly.com/content/generic-mitigations



Generic mitigations

oreilly.com/content/generic-mitigations

Binary rollback

Data rollback

Degrade

Upsize

Block/Quarantine

Drain



Dynamically moving workloads



Reliability cannot be an afterthought



So *what* did we *do*?



Closing thoughts



Reliability cannot be an afterthought
...if you're building reliable systems



