

Scaling Configuration Management Across Data Center And Cloud

June 12, 2012

Ted Zlatanov

Senior CFEngine Architect

CFEngine AS



ted.zlatanov@cfengine.com



[@cfengine_news](https://twitter.com/cfengine_news)



Introduction

- Configuration Management: a fun field with great companies, great tools: CFEngine, PuppetLabs, Opsware, bcfg2, and others
- Common needs at any scale: inventory management, configuration management, ease of use



Special Needs

- Enterprise needs: reporting, compliance, performance
- Cloud needs: tracking instances, scaling from 1 to N, forget the hype
- Networking needs: provisioning, IPv6, provisioning and monitoring ports and protocols
- http://blogs.forrester.com/james_staten/12-05-24-cloud_inefficiency_bad_habits_are_hard_to_break



Infrastructure Needs for Scaling

- Infrastructure discovery: Zeroconf/DNS-SD, multicast, simple UDP, traditional (often painfully manual) registration methods
- Inventory management: simple, fast, portable, flexible, rule-based or hierarchy-based
- Instance provisioning and management, tracking, and monitoring



CM-related Infrastructure Needs for Scaling

- Performance and efficiency across the machine's lifecycle
- Agent reporting across 1 to N instances, and efficient aggregation
- Efficient distribution of policies and metadata
- Knowledge is key at the top level, and policies are not static things, but manipulations and adaptations of knowledge into action



CAP Theorem for CM

- Originally, CAP referred to databases (Brewer, “Towards robust distributed systems”, PoDC 2000)
- It’s not just “pick 2 out of 3: consistency, availability, partitioning tolerance”!
- Further discussed by Gilbert and Lynch, and Brewer elaborated recently in <http://www.infoq.com/articles/cap-twelve-years-later-how-the-rules-have-changed>



CAP Theorem for CM

- Insight: consistency is a state, but CM convergence is a process, so CM is always balancing. The original CAP theorem lacks the time element, especially latency, so we believe convergence is simply the extension of consistency along time.
- Research is ongoing, talk to us if you want to know more



Collaboration! Standards!

- We need to collaborate to grow the CM space!!!
- Benchmarks: we need neutral benchmarks of key characteristics of CM systems, especially performance and resource usage in cloud environments with constrained instances



Collaboration! Standards!

- Inventory Management: share what you know. Make it easy to move between CM systems and use more than one at a time. Locking customers into a single CM system is SO last millennium.
- Common protocols: let's establish ways to communicate what happened, who did it, and why, in a neutral way.



Thank you

Ted Zlatanov

Senior CFEngine Architect

CFEngine AS

ted.zlatanov@cfengine.com

 [@cfengine_news](https://twitter.com/cfengine_news)

