

Thermal Load Balancing

PlateSpin Orchestrate use case

Adam Spiers

aspiers@novell.com

Novell[®]

The Problem: Data Center Cooling

- Cooling is expensive
 - In a typical data center, more power is spent on cooling than on servers. Estimates range from 44%¹ to as high as 63%.²
- Cooling is unavoidable
 - 10% of racks have ambient temperatures of 75°F or higher at the air intake at the top of the rack. High temperatures are causes of decreased hardware reliability. Intermittent ghosts and outright hardware failures are three times more prevalent in the top third of racks than the bottom two-thirds.³
- Air conditioning units are critical to service availability
 - Servers can redline within 90 seconds of an AC unit failure.⁴

¹ See, Simon. *Is there a pathway to a Green Grid?* Sun, 2008 <http://www.ibergrid.eu/2008/presentations/Dia%2013/4.pdf>

² <http://h20331.www2.hp.com/ERC/cache/438048-0-0-225-121.html>

³ Sullivan, Robert F., Ph.D. *Reducing Bypass Airflow Is Essential for Eliminating Hotspots*. Upsite Technologies <http://www.42u.com/data-center-hot-spots.htm>

⁴ Sharma, Ratnesh et al. *Balance of Power: Dynamic Thermal Management for Internet Data Centers*. HP Labs, 2003 <http://www.hpl.hp.com/techreports/2003/HPL-2003-5.pdf>

The Problem: Data Center Cooling (continued)


- Many data centers cool incorrectly
 - Cooling over-capacity is very common, and is not a predictor of successful cooling. In one study, nineteen rooms studied ran on average 2.7 times more cooling equipment than required to cool the computer heat load. Two rooms ran **16 times** more cooling than required, yet one had 20% hot racks/cabinets and the other had 7% hot racks/cabinets.¹
- There are over 80 energy efficiency incentive or rebate programs offered by local utilities or state energy efficiency programs in the US alone.²

¹ Sullivan, Robert F., Ph.D. *Reducing Bypass Airflow Is Essential for Eliminating Hotspots*. Upsite Technologies <http://www.42u.com/data-center-hot-spots.htm>

² *The green data center*. IBM, May 2007 http://www-900.ibm.com/cn/systems/migratetoibm/pdf/Energy-file03_OIW03002USEN.pdf

Existing solutions are limited

- Traditional raised floor cooling with alternating hot/cold aisles can typically only handle up to 5 kW per rack¹
 - but rack power consumption can be in excess of 10kW
 - and the evolution of blade designs are causing density and power per rack to continually increase



©2004 Liebert Corporation. All Rights Reserved.

¹ See, Simon. *Is there a pathway to a Green Grid?* Sun, 2008 <http://www.ibergrid.eu/2008/presentations/Dia%2013/4.pdf>

Existing solutions are costly

- Other solutions involve expensive, highly specialised thermal management, e.g.
 - pumped refrigerant coolant systems
 - hot air containment
 - > per rack / per aisle
 - > per data center! (“DC in a box”)
 - variable rate cooling units
 - sub-floor air pressure management
- Are there cheaper approaches?
- Can we **treat the cause** rather than the symptoms?
- What if we could do something *in software*?

A Novel (Novell®) Solution: Dynamic Thermal Load Balancing

- Virtualization and live migration allows dynamic relocation of workloads with no impact on service
 - Policy-based migration of VMs from hot spots to cool spots
- Energy consumption can be reduced by more than 14% by intelligent workload placement.¹
- Risk of service outage due to computer room air conditioning unit failures can be mitigated by migrating workloads away from the failed unit.¹
- Automatic VM live migration based on thermal policies can be implemented easily using PlateSpin® Orchestrate from Novell.

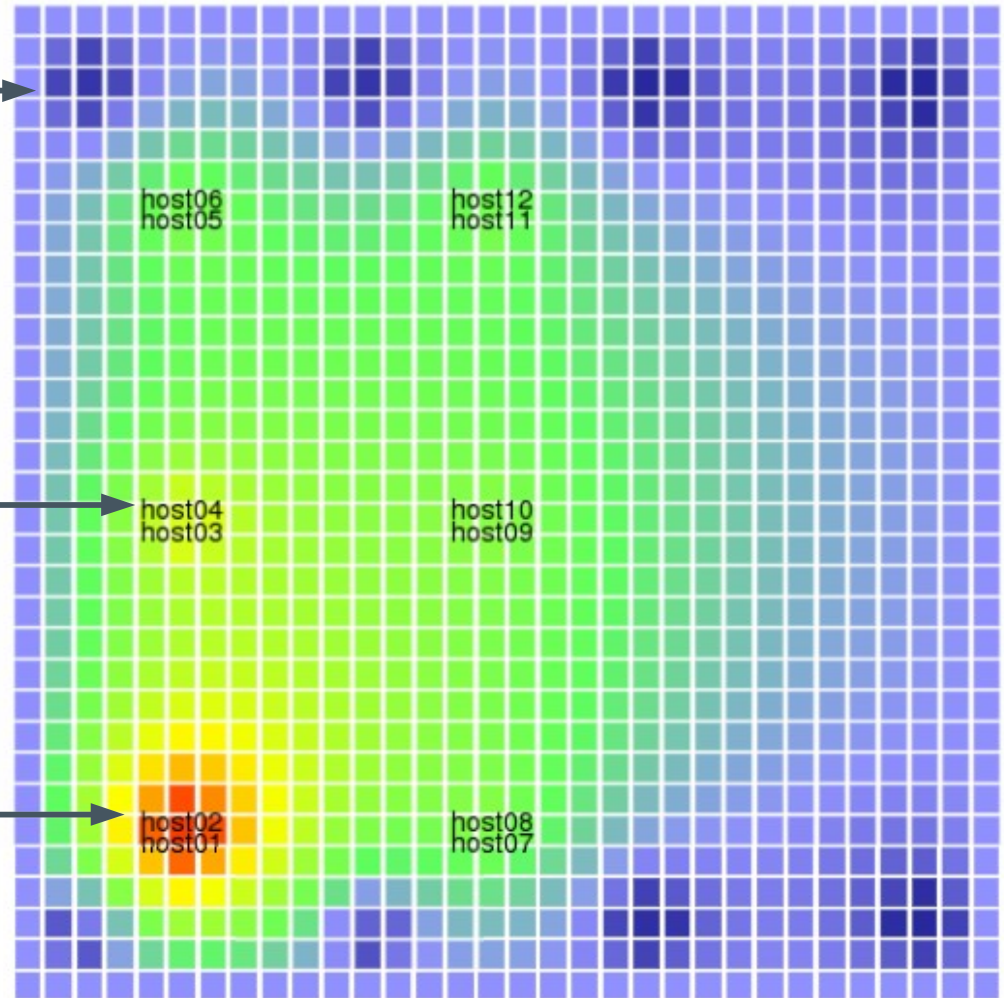
¹ Sharma, Ratnesh et al. *Balance of Power: Dynamic Thermal Management for Internet Data Centers*. HP Labs, 2003
<http://www.hpl.hp.com/techreports/2003/HPL-2003-5.pdf>

Thermal Load Balancing

Cooling Units

“Rack” of VM host hardware operating below capacity

VM host hardware operating at or near capacity



Novell®

Unpublished Work of Novell, Inc. All Rights Reserved.

This work is an unpublished work and contains confidential, proprietary, and trade secret information of Novell, Inc. Access to this work is restricted to Novell employees who have a need to know to perform tasks within the scope of their assignments. No part of this work may be practiced, performed, copied, distributed, revised, modified, translated, abridged, condensed, expanded, collected, or adapted without the prior written consent of Novell, Inc. Any use or exploitation of this work without authorization could subject the perpetrator to criminal and civil liability.

General Disclaimer

This document is not to be construed as a promise by any participating company to develop, deliver, or market a product. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. Novell, Inc. makes no representations or warranties with respect to the contents of this document, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. The development, release, and timing of features or functionality described for Novell products remains at the sole discretion of Novell. Further, Novell, Inc. reserves the right to revise this document and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes. All Novell marks referenced in this presentation are trademarks or registered trademarks of Novell, Inc. in the United States and other countries. All third-party trademarks are the property of their respective owners.

