CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

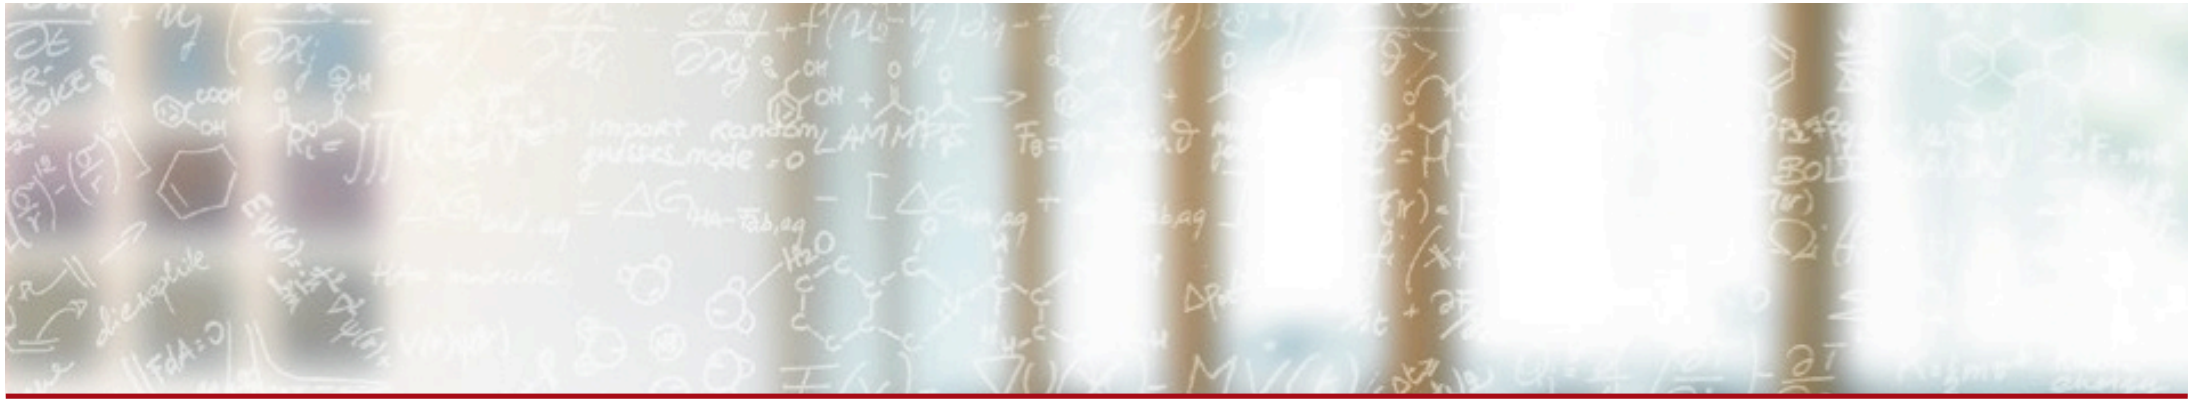

Incorporating a Test and Development System Within the Production System

Cray User Group 2018
Nicholas P. Cardo, CSCS
Marco Induni, CSCS
May DD, 2018

Outline

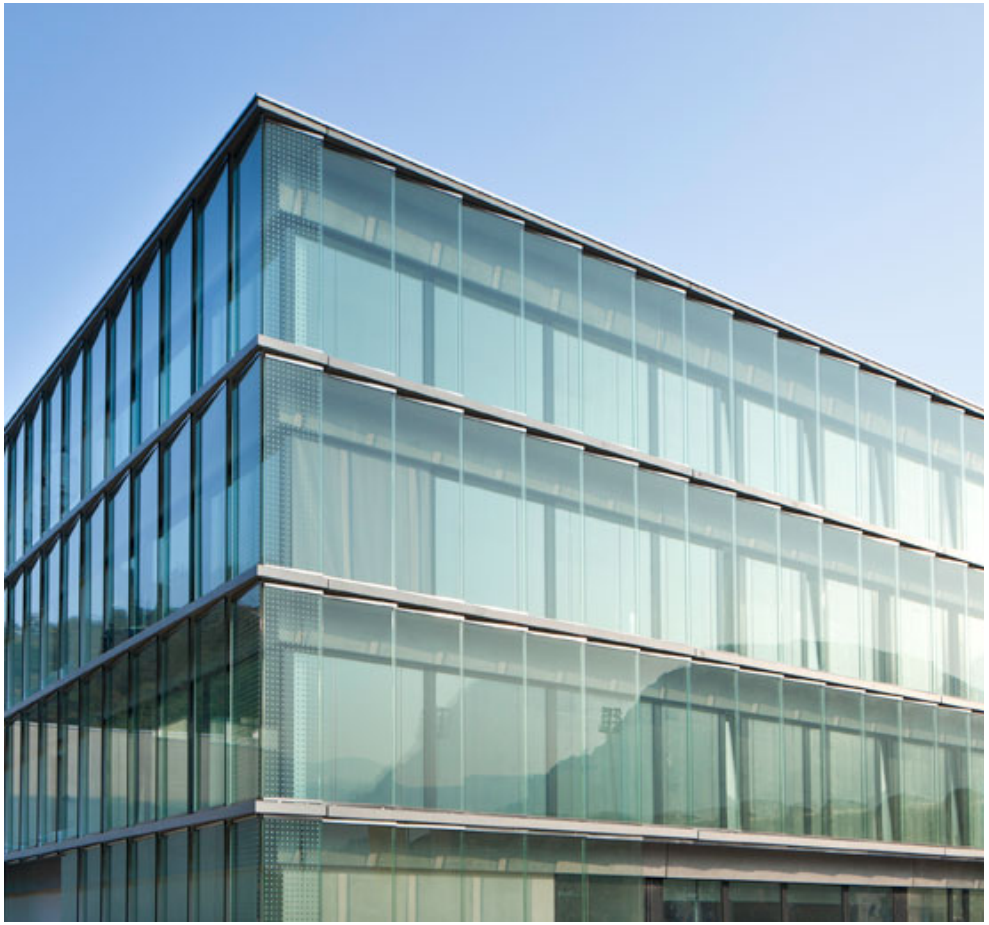

- Test and Development Systems
- The Problem...
- Systems Description
- Motivation
- Implementation
- Challenges
- Story of Success

Value of a Test and Development System

- Evaluate the impact of new software levels
  - Without impacting production operations
- Provide an upgraded environment to rebuild and test applications
  - Mission Critical applications MUST work on new software/OS levels
  - Rebuilding large applications can take a significant effort and time
- Provide an environment for experimentation
  - Develop and optimize processes and procedures

The Problem

- So much money, so little hardware…
- Tradeoff
  - Buy more compute capacity
  - Buy TDS capability
- In Production HPC, TDS should prevail
  - But often doesn't…

https://www.1843magazine.com/content/ideas/ian-leslie/non-cogito-ergo-sum

Hardware Description

- **Compute Nodes (12)**
  - 8 x NVIDIA Tesla K80 GPUs
  - 256 GB of memory
  - 2 x Intel Xeon CPU E5-2690 v3
- **Post Processing Nodes (5)**
  - 256 GB of memory
  - 2 x Intel Xeon CPU E5-2690 v3
- **Login Nodes (3)**
  - 128 GB of memory
  - 2 x Intel Xeon CPU E5-2690 v3

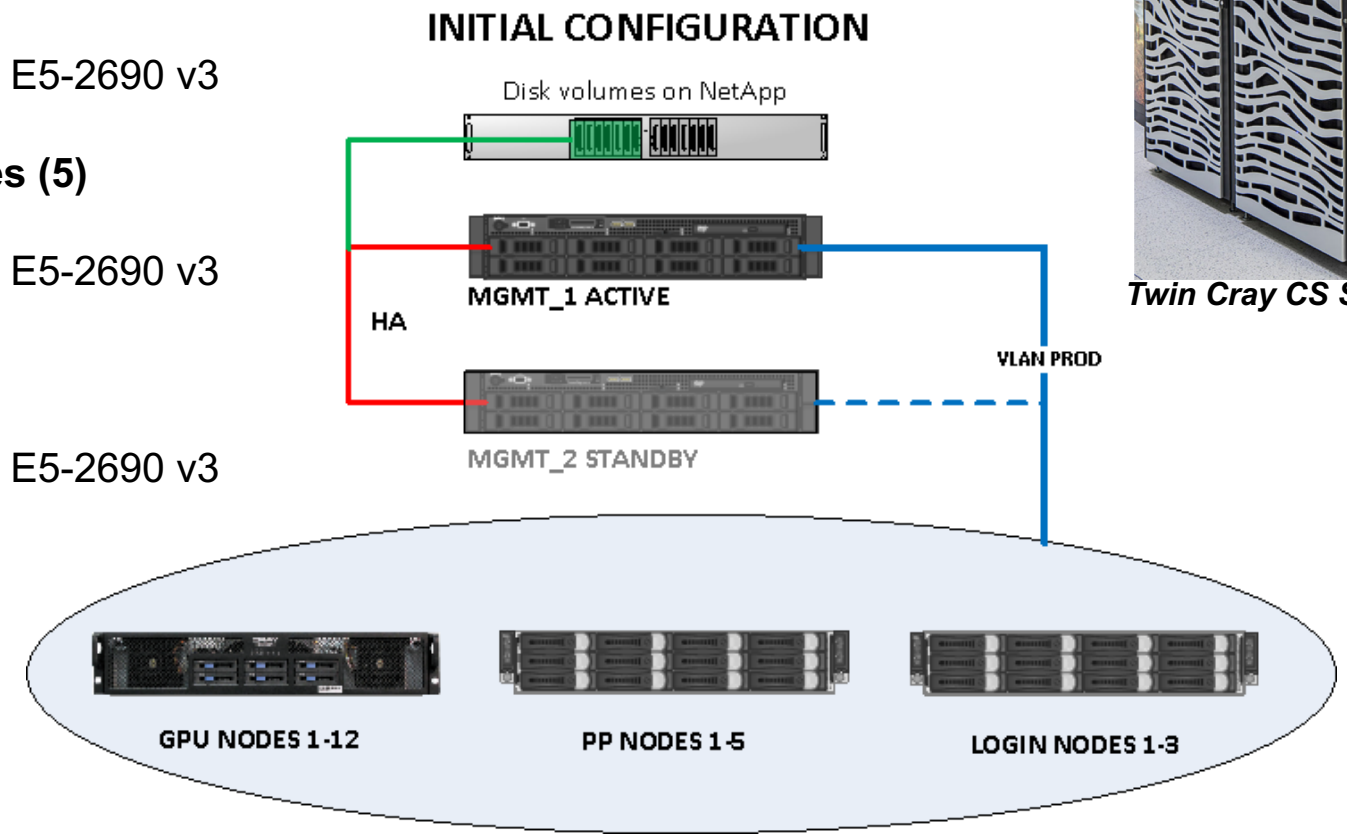

Twin Cray CS Storm Systems

Motivation

- Production weather forecasting system for Switzerland
  - Possibilities for test time are few and far between and short
- Without a TDS
  - Higher risk of introducing problems during tests
- Testing is time consuming
  - Many man-hours required to rebuild and re-validate
- Major upgrades are disruptive
  - Red Hat 6 -> Red Hat 7
- Specialized hardware is very expensive
  - And no money…

http://www.businessmantraa.in/short-quotes-explanation-motivation-success-business/

The Big What If…

What if we could
“borrow” hardware from
the production system
and use it like a TDS
and pigs could fly?

https://healthfacts.ng/positive-thinking-and-your-health-1/

http://kineticmotions.ca/what-if

This simple question set in motion a series of events that would ultimately solve the problem.

Implemention Details

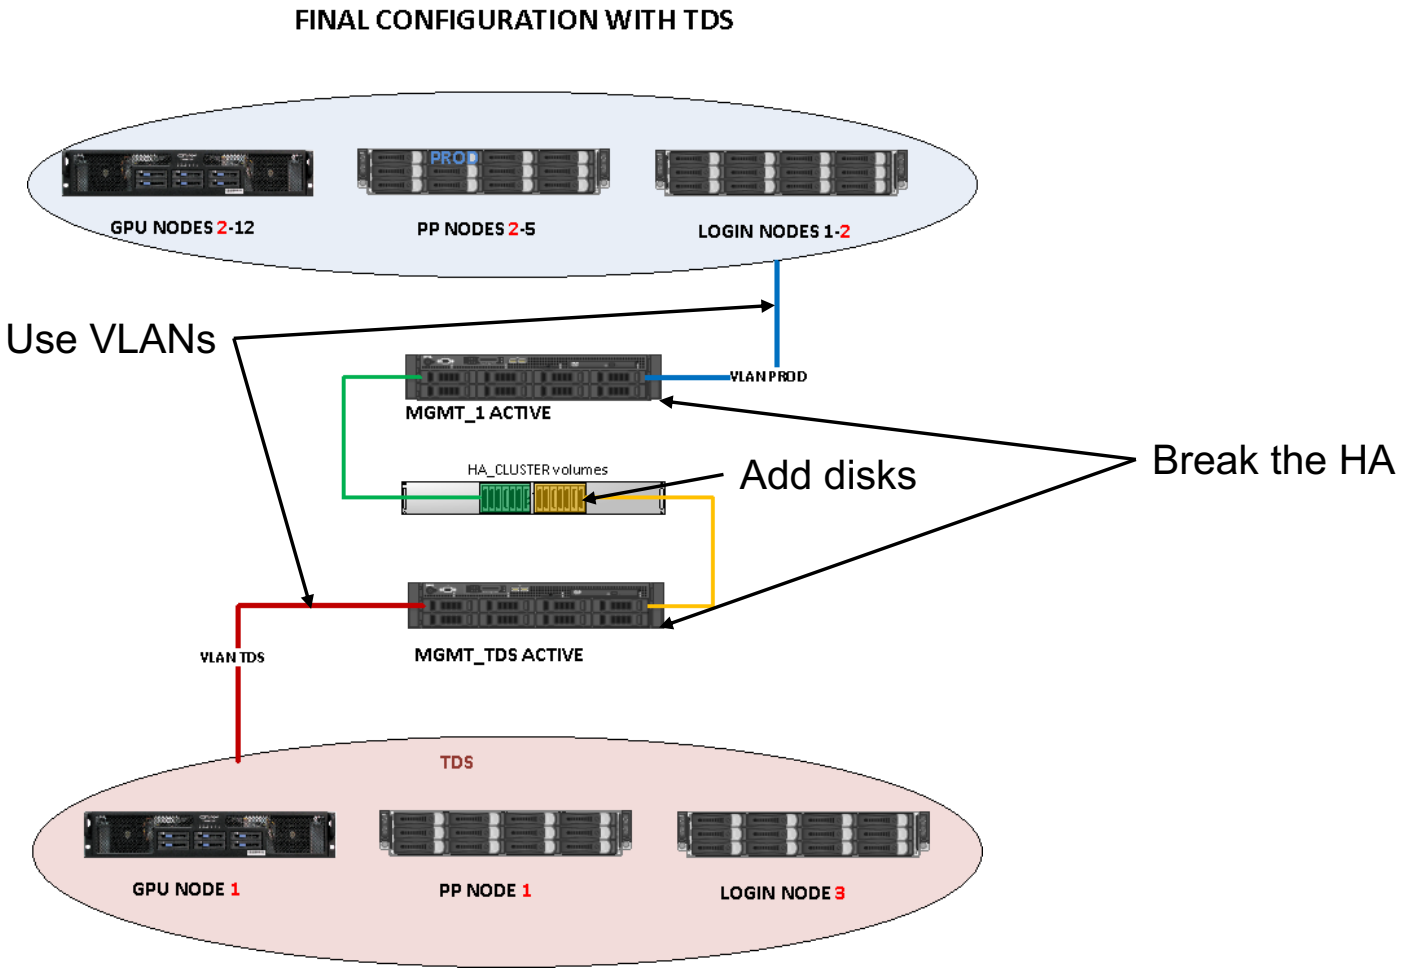

Challenges

- No upgrade path to Red Hat 7 for the Cray Advanced Cluster Engine (ACE)
  - After 7 months, enough was enough
    - Switched to Bright Cluster Manager (BCM), through Cray
- No migration path from ACE to BCM
  - Need to do a fresh install
  - Took a while to get all the rpms right
- No High Availability for the system if H/A is broken
  - Acceptable risk due to other redundancies
- How many custom images?
  - End result only 1, personalize at boot time

And the Second System?

- Process repeated to prepare second system
  - Much faster having already worked it out
- Complicated but careful migration plan to new O/S
  1. Boot entire backup system to new O/S
  2. Switch production to backup system
  3. Run for 24 hours
  4. Boot primary system to new O/S
  5. Switch production to primary system
- Safety Precautions
  - Previous O/S available on second management workstations
  - A reboot returns system to previous levels

A Success Story

- Very Successful!
  - No interruption to production runs
  - Viable solution for future software updates
    - Capability remains in place
    - Red Hat 7.3 -> Red Hat 7.5 later this year…
- But still, there are challenges
  - Restoring H/A to management workstations

It Takes People…

- Thank-You to Cray for helping to work through our problems and frustrations
- Thank-You to Bright Computing for helping with the migration to BCM

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

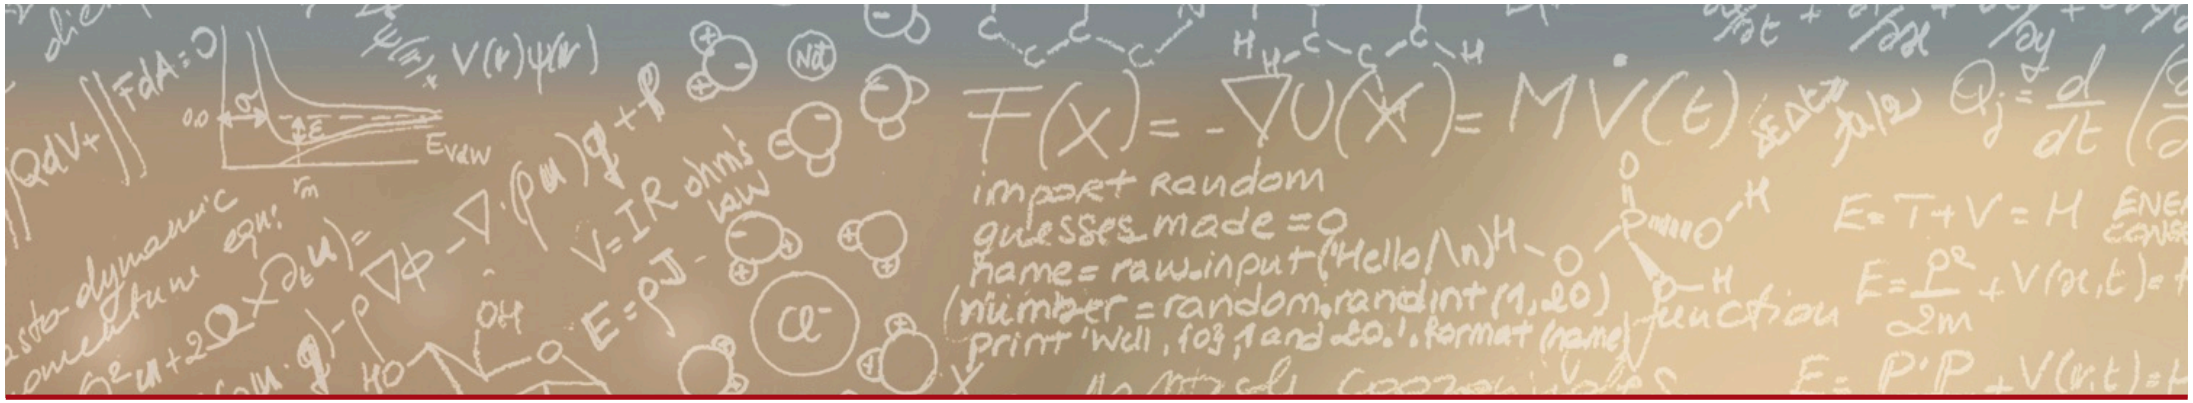

Grazie per la vostra attenzione.
Thank-You for your attention.