QuantCatalyst

A Cluster is More Than Just a Bunch of Boxes

Tilman Koschnick
September 2008

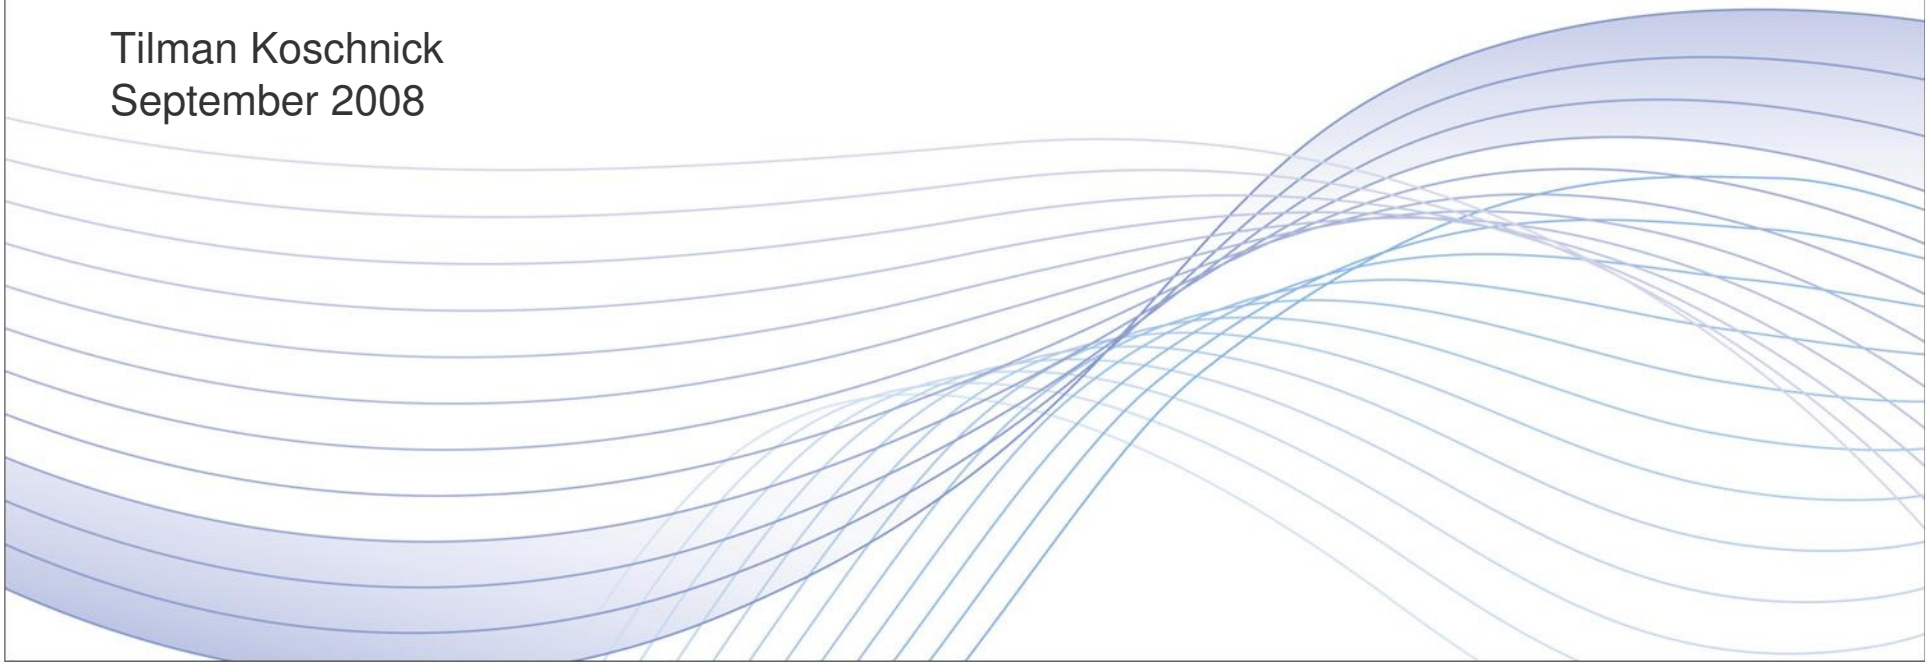

Contents

Contents

A “Cluster” is a bunch of boxes…

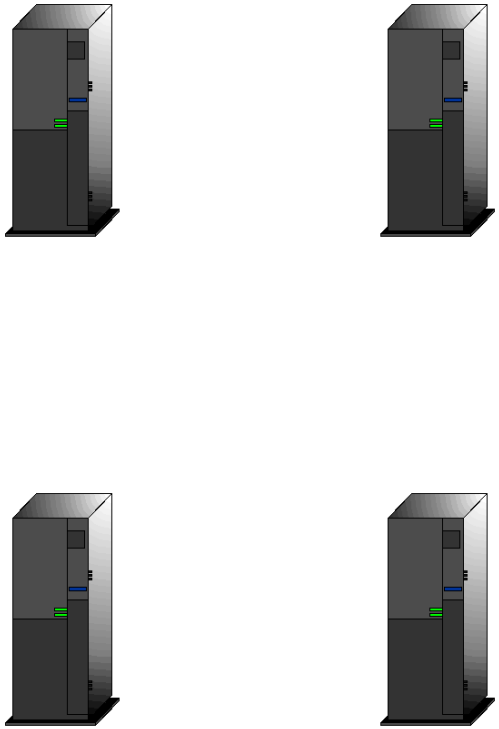

Quant Catalyst

… connected by a network

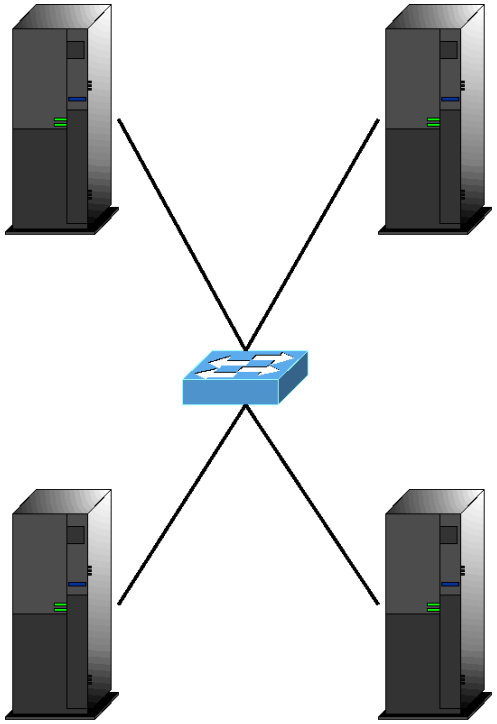

High Performance Computing cluster

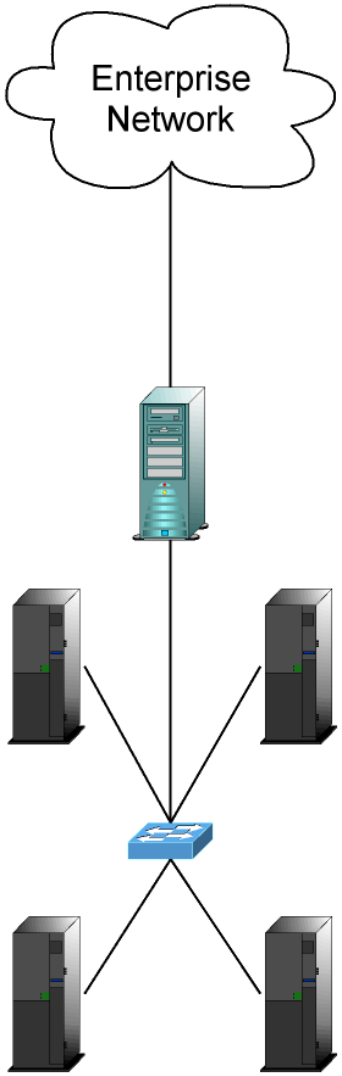

Quant Catalyst

High Availability cluster

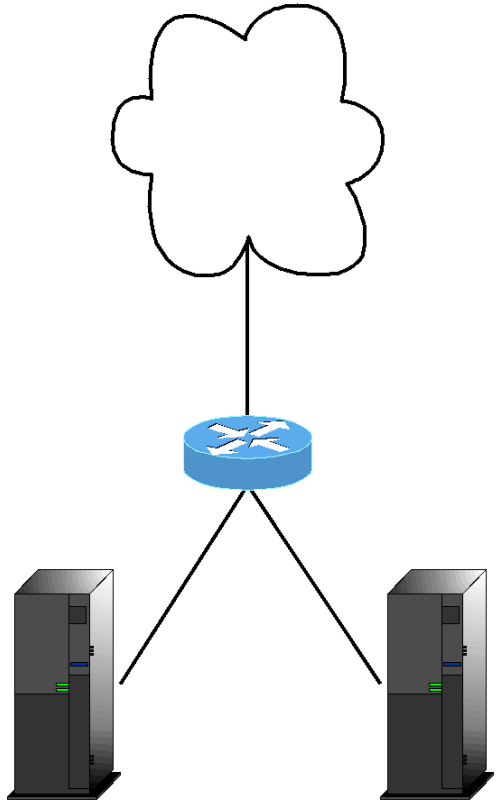

Quant Catalyst

High Availability cluster

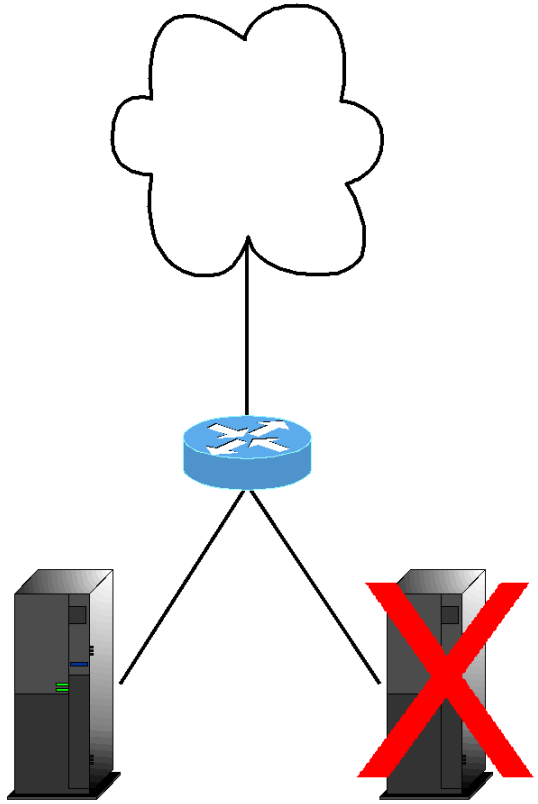

Database cluster

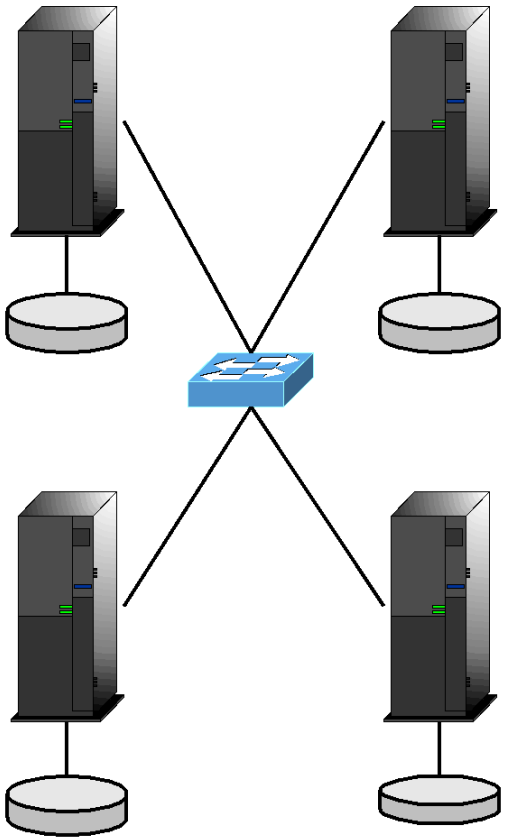

Grid of clusters

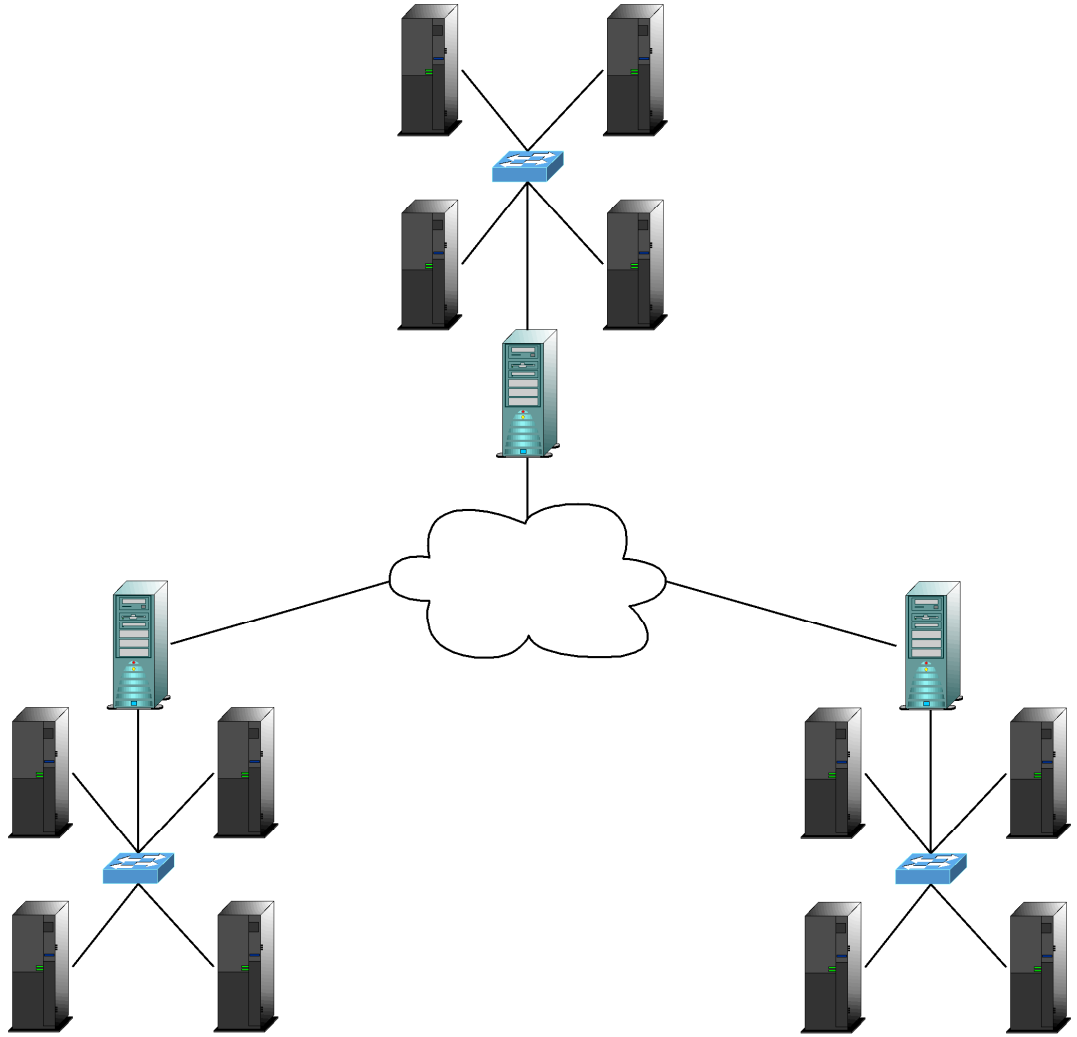

Grid of workstations

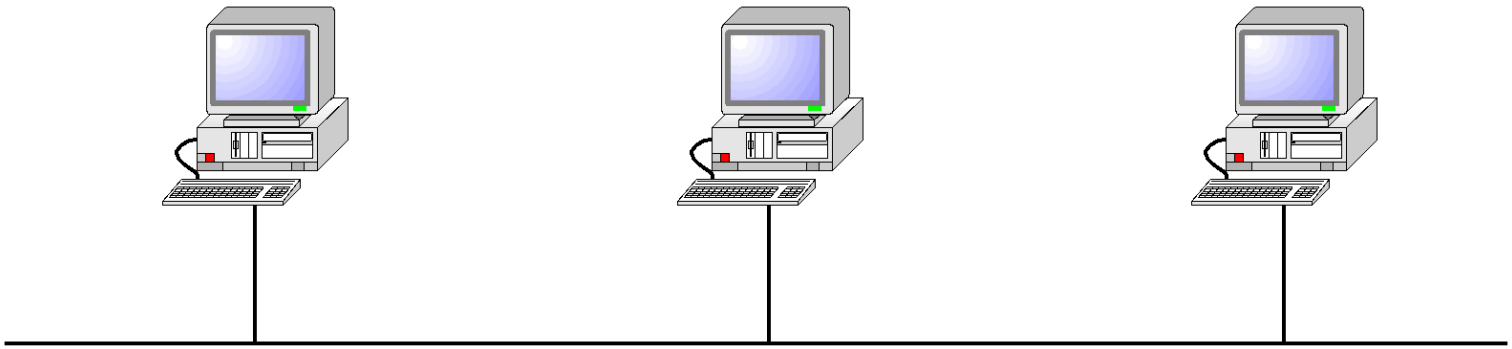

Quant Catalyst

Definitions

- Cluster
  - tightly connected
  - High Performance Computing (HPC) cluster
  - High Availability (HA) cluster
  - Database (DBMS) cluster
- Grid
  - loosely connected
  - grid of clusters
  - grid of workstations

Contents

History

- First HPC cluster:
  - “Beowulf”
  - 1993/1994 at NASA
  - built on top of emerging Open Source software / Linux
- Prevailing technology at the time:
  - multi-processor & vector supercomputer
  - highly specialised, purpose build and expensive
- The Beowulf approach:
  - economy of scale
  - use commodity off the shelf (COTS) components
  - history repeats: GPU from gamer’s PCs, cell processor from playstation

Cluster adoption in HPC

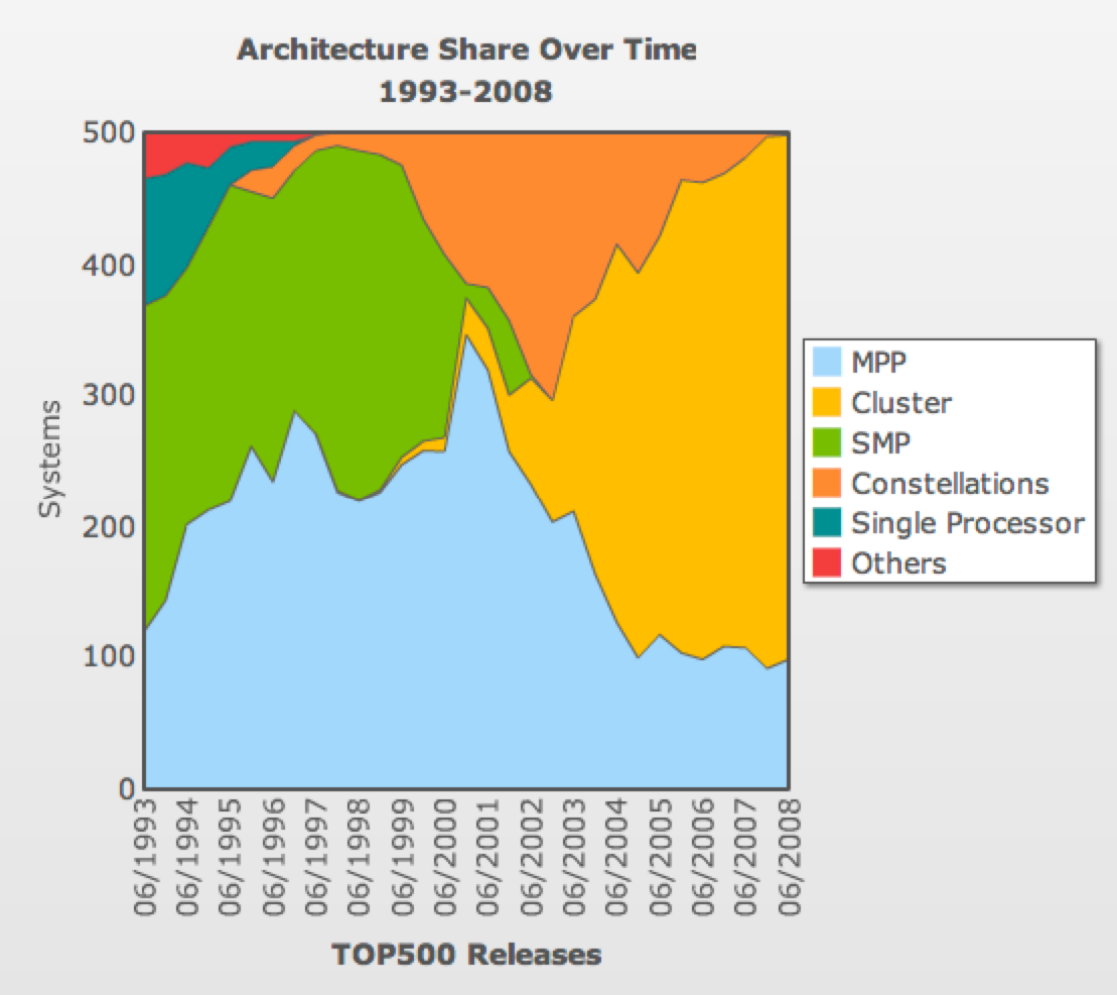

Source: http://www.top500.org/overtime/list/31/archtype

Quant Catalyst

Contents

Challenges

- Similar to other enterprise networks
- Easier: more homogeneous
- Harder: more homogeneous

Challenges

- Two examples:

- Keep large number of machines in sync
  - roll out software and configuration across many machines at once
  - deal with downtime of single nodes during roll out

- Security
  - protect many machines
  - balance performance and security

Contents

Components of Cluster Management

- Package (or software) management
- Configuration management
- Monitoring

Package management

- Package management provides means to:
  - distribute
  - install
  - upgrade
  - remove
- … and ensures full accountability:
  - what
  - when
  - where
  - which version
- Install **all** software through package management:
  - system packages
  - 3rd party packages
  - applications

Quant Catalyst

Configuration management

- Configuration is the sum of all:
  - system parameters
  - software settings
- Configuration manager:
  - runs on every machine in the cluster
  - monitors and adjusts settings and parameters
- Configuration is based on classes and rules
  - rules describe what to do and how to do it
  - classes describe where to apply rules

Configuration management

- Rules describe the target state
  - configuration manager compares with current state
  - activates rule if current and target state do not match
- Rules are generally simple:
  - for a given class:
  - copy config file from configuration repository…
  - … then restart associated service
- Rules and config files are kept under version control
  - central repository with configuration information
  - preserves history, provides documentation

Monitoring

- Fault monitoring
  - report irregularities
  - watch for failed or soon-to-fail hardware and software components
  - …
- Resource monitoring
  - free disk space
  - processor and memory usage
  - network utilisation
  - …
- Performance monitoring
  - job duration
  - job throughput
  - overall cluster utilisation
  - …

Contents

Principles of Cluster Management

- Automate
- Centralise
- Unify

Automate

- Don’t repeat yourself (DRY)
- Automation is documentation
  - rules and scripts describe how to get from A to B
- … is reproducible
  - scripts deliver consistent results, any time, on any machine
- … is humane
  - it is more fun to teach a computer to do repetitive tasks than to do them yourself
  - repetitive, boring tasks lead to errors and inconsistencies

Centralise

- Treat a cluster **not** as a bunch of machines, but as a single system
  - the complexity of a cluster should be hidden from users
  - … and system administrators as well, wherever possible!
- Provide/require only one interface to the enterprise network
  - provide single system view for integration tasks
- If you have to tweak something, tweak it once, in one place only
  - distribute all changes from one single repository
- Collect all information in one place
  - collect all logging and monitoring information in one single place

Centralise

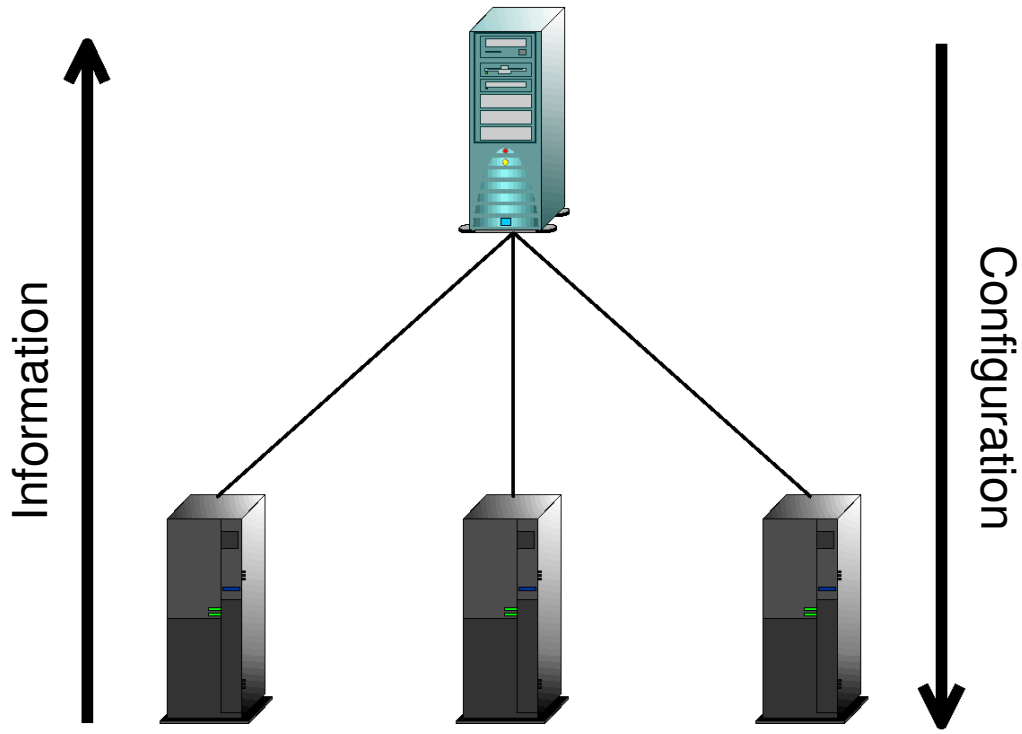

Quant Catalyst

Centralise – Scaling and Redundancy

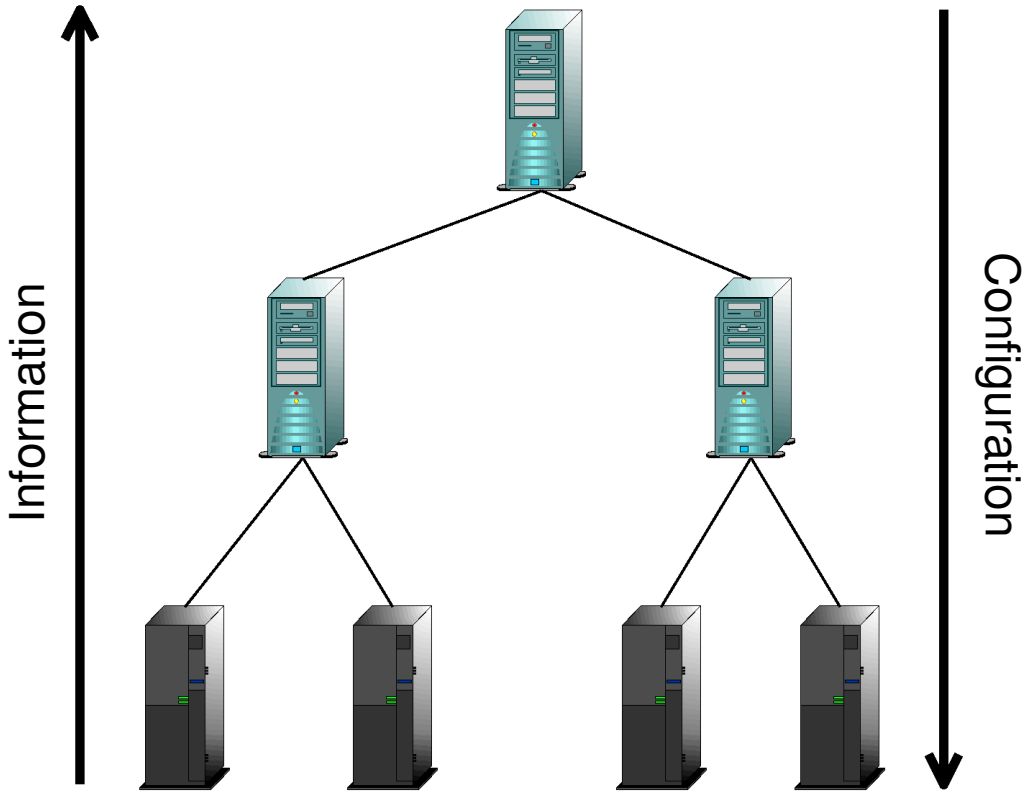

Quant Catalyst

Unify

- One mechanism to:
  - install, update and maintain the system
  - recover from disaster
- On installation
  - bootstrap with a minimal system image, installed over the network
  - afterwards, package & configuration management takes over
- Update & maintenance
  - fully employ package & configuration management
- Disaster recovery
  - the same as installation
  - disaster recovery strategies are continually kept up to date and exercised

Contents

Problem solved: Synchronisation

- Problems:
  - synchronised and identical roll out
  - node failure during roll out

- Solution:
  - all nodes installed and configured automatically and identically
  - automatic updates as soon as failed node is back online

Problem solved: Security

- Problems:
  - keep large number of computers secure
  - balance performance and security
- Solution:
  - cluster seen as one security domain
  - expensive security measures within domain not required
  - single controlled gateway to enterprise network
  - secure configuration enforced…
  - … controlled and supervised

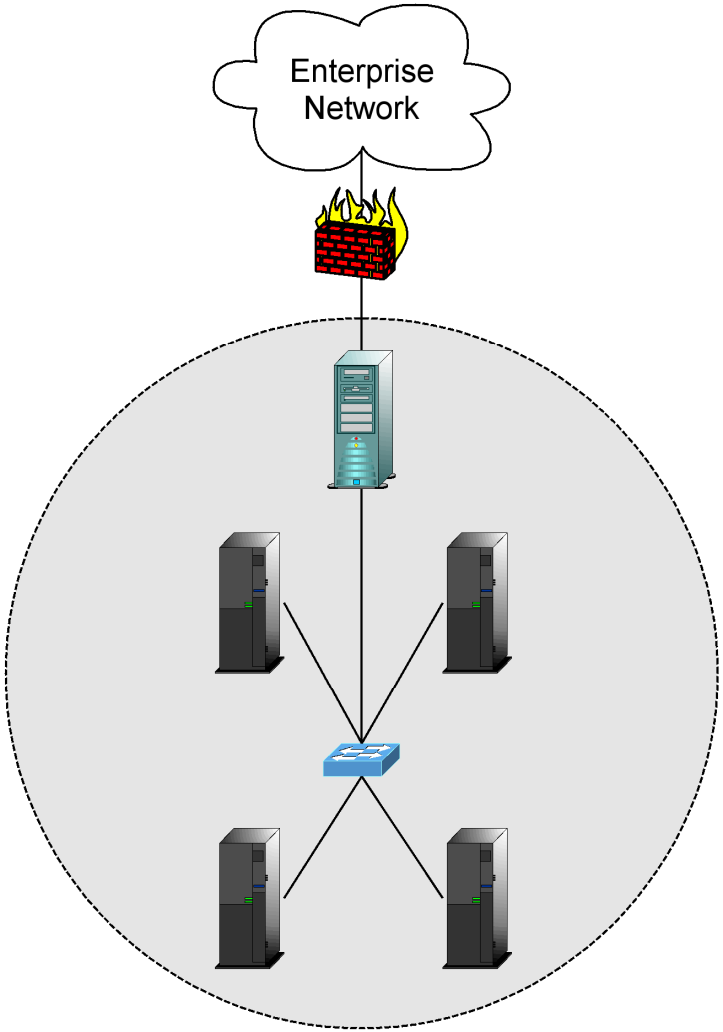

Quant Catalyst

Conclusions

- It takes some effort and some care to turn a bunch of boxes into a smoothly running cluster.
- The use of some core components and the application of a couple of central principles goes a long way.
- Questions?

Quant Catalyst

Contact

Tilman Koschnick
tilman.koschnick@quantcatalyst.com
Phone: +41 43 477 53 30
Mobile: +41 76 592 12 51

Dr Daniel Egloff
daniel.egloff@quantcatalyst.com
Phone: +41 44 520 01 17
Mobile: +41 79 430 03 61

Quant Catalyst

Hardturmstrasse 101
CH-8005 Zürich