LHCb: experience, feedback and requirements

M. Clemencic

CERN - LHCb

June 21, 2010
2nd Workshop on adapting applications and computing services to multi-core and virtualization

Outline

Introduction

- HEP has always been good at using computing resources
  - natural and easy parallelization model
- Technology evolution
  - Single-Core $\rightarrow$ Multi-Core $\rightarrow$ Many-Core
  - Computers $\rightarrow$ Virtual machines $\rightarrow$ Computing clouds
- New models needed

Introduction

- HEP has always been good at using computing resources
  - natural and easy parallelization model
- Technology evolution
  - Single-Core $\rightarrow$ Multi-Core $\rightarrow$ Many-Core
  - Computers $\rightarrow$ Virtual machines $\rightarrow$ Computing clouds
- New models needed

Introduction

- HEP has always been good at using computing resources
  - natural and easy parallelization model
- Technology evolution
  - Single-Core $\rightarrow$ Multi-Core $\rightarrow$ Many-Core
  - Computers $\rightarrow$ Virtual machines $\rightarrow$ Computing clouds
- New models needed

Introduction

- HEP has always been good at using computing resources
  - natural and easy parallelization model
- Technology evolution
  - Single-Core $\rightarrow$ Multi-Core $\rightarrow$ Many-Core
  - Computers $\rightarrow$ Virtual machines $\rightarrow$ Computing clouds
- New models needed

Introduction

- HEP has always been good at using computing resources
  - natural and easy parallelization model
- Technology evolution
  - Single-Core $\rightarrow$ Multi-Core $\rightarrow$ Many-Core
  - Computers $\rightarrow$ Virtual machines $\rightarrow$ Computing clouds
- New models needed

Introduction

- HEP has always been good at using computing resources
  - natural and easy parallelization model
- Technology evolution
  - Single-Core $\rightarrow$ Multi-Core $\rightarrow$ Many-Core
  - Computers $\rightarrow$ Virtual machines $\rightarrow$ Computing clouds
- New models needed

Introduction

- HEP has always been good at using computing resources
  - natural and easy parallelization model
- Technology evolution
  - Single-Core $\rightarrow$ Multi-Core $\rightarrow$ Many-Core
  - Computers $\rightarrow$ Virtual machines $\rightarrow$ Computing clouds
- New models needed

Introduction

- HEP has always been good at using computing resources
  - natural and easy parallelization model
- Technology evolution
  - Single-Core $\rightarrow$ Multi-Core $\rightarrow$ Many-Core
  - Computers $\rightarrow$ Virtual machines $\rightarrow$ Computing clouds
- New models needed

Introduction

- HEP has always been good at using computing resources
  - natural and easy parallelization model
- Technology evolution
  - Single-Core $\rightarrow$ Multi-Core $\rightarrow$ Many-Core
  - Computers $\rightarrow$ Virtual machines $\rightarrow$ Computing clouds
- New models needed

Outline

The market

- Technologies
  - Virtualization
    - Xen, KVM, VMWare, Parallels, VirtualBox, Hyper-V, ...
    - Virtual machine $\rightarrow$ Virtual computing centers
  - Operating systems
    - Appliances: CernVM, TurnKey, ...
    - Regular OSs: Ubuntu, RedHat, Scientific Linux, ...
- Services
  - Virtual computing centers, clouds

The market

- Technologies
  - Virtualization
    - Xen, KVM, VMWare, Parallels, VirtualBox, Hyper-V, ...
    - Virtual machine $\rightarrow$ Virtual computing centers
  - Operating systems
    - Appliances: CernVM, TurnKey, ...
    - Regular OSs: Ubuntu, RedHat, Scientific Linux, ...
- Services
  - Virtual computing centers, clouds

The market

- Technologies
  - Virtualization
    - Xen, KVM, VMWare, Parallels, VirtualBox, Hyper-V, ...
    - Virtual machine $\rightarrow$ Virtual computing centers
  - Operating systems
    - Appliances: CernVM, TurnKey, ...
    - Regular OSs: Ubuntu, RedHat, Scientific Linux, ...
- Services
  - Virtual computing centers, clouds

Features

- Isolation from the host system
- General users
  - easy installation
  - minimal impact on the host
  - run on unsupported platforms
- Computing centers
  - system configuration tuned for the users
  - partitioning of computing farms
  - security

Features

- Isolation from the host system
- General users
  - easy installation
  - minimal impact on the host
  - run on unsupported platforms
- Computing centers
  - system configuration tuned for the users
  - partitioning of computing farms
  - security

Features

- Isolation from the host system
- General users
  - easy installation
  - minimal impact on the host
  - run on unsupported platforms
- Computing centers
  - system configuration tuned for the users
  - partitioning of computing farms
  - security

Features

- Isolation from the host system
- General users
  - easy installation
  - minimal impact on the host
  - run on unsupported platforms
- Computing centers
  - system configuration tuned for the users
  - partitioning of computing farms
  - security

Special case: LHCb Online

Microsoft Hyper-V being used to redefine the service infrastructure (web servers, firewalls, gateways, . . .), with plans to (re)evaluate KVM in REL6.

Features

- Isolation from the host system
- General users
  - easy installation
  - minimal impact on the host
  - run on unsupported platforms
- Computing centers
  - system configuration tuned for the users
  - partitioning of computing farms
  - security

Special case: LHCb Online

Microsoft Hyper-V being used to redefine the service infrastructure (web servers, firewalls, gateways, . . .), with plans to (re)evaluate KVM in REL6.

Outline

How can we use it?

- Abstraction from the host to simplify portability
- CernVM
  - Centrally managed configuration
  - Reduced disk and bandwidth requirements
  - Network configuration issues
- Other systems
  - Compatibility issues
- Basic requirements
  - Standard service
  - Supported by the batch systems
  - Transparent for the users

How can we use it?

- Abstraction from the host to simplify portability
- CernVM
  - Centrally managed configuration
  - Reduced disk and bandwidth requirements
  - Network configuration issues
- Other systems
  - Compatibility issues
- Basic requirements
  - Standard service
  - Supported by the batch systems
  - Transparent for the users

How can we use it?

- Abstraction from the host to simplify portability
- CernVM
  - Centrally managed configuration
  - Reduced disk and bandwidth requirements
  - Network configuration issues
- Other systems
  - Compatibility issues
- Basic requirements
  - Standard service
  - Supported by the batch systems
  - Transparent for the users

How can we use it?

- Abstraction from the host to simplify portability
- CernVM
  - Centrally managed configuration
  - Reduced disk and bandwidth requirements
  - Network configuration issues
- Other systems
  - Compatibility issues
- Basic requirements
  - Standard service
  - Supported by the batch systems
  - Transparent for the users

How can we use it?

- Abstraction from the host to simplify portability
- CernVM
  - Centrally managed configuration
  - Reduced disk and bandwidth requirements
  - Network configuration issues
- Other systems
  - Compatibility issues
- Basic requirements
  - Standard service
  - Supported by the batch systems
  - Transparent for the users

Outline

Why

- Evolution of CPUs
  - past: cores with more and more power
  - future: more and more cores with less power
  - GPUs can be used for computation
- Our parallelization model will soon stop scaling
  - RAM technology doesn't evolve as quick as CPUs
  - GPUs cannot run full applications
- We need to take actions to be able to squeeze all the FLOPS out of the machines

Why

- Evolution of CPUs
  - past: cores with more and more power
  - future: more and more cores with less power
  - GPUs can be used for computation
- Our parallelization model will soon stop scaling
  - RAM technology doesn't evolve as quick as CPUs
  - GPUs cannot run full applications
- We need to take actions to be able to squeeze all the FLOPS out of the machines

Why

- Evolution of CPUs
  - past: cores with more and more power
  - future: more and more cores with less power
  - GPUs can be used for computation
- Our parallelization model will soon stop scaling
  - RAM technology doesn't evolve as quick as CPUs
  - GPUs cannot run full applications
- We need to take actions to be able to squeeze all the FLOPS out of the machines

Why

- Evolution of CPUs
  - past: cores with more and more power
  - future: more and more cores with less power
  - GPUs can be used for computation
- Our parallelization model will soon stop scaling
  - RAM technology doesn't evolve as quick as CPUs
  - GPUs cannot run full applications
- We need to take actions to be able to squeeze all the FLOPS out of the machines

How

- Old technologies
  - Multi-processing
  - Multi-threading
- New technologies
  - OpenMP
  - Threading Building Blocks
  - Concurrent Collections
  - ...
- New approaches
  - Parallelization-friendly languages (Haskell, Erlang, ...)

How

- Old technologies
  - Multi-processing
  - Multi-threading
- New technologies
  - OpenMP
  - Threading Building Blocks
  - Concurrent Collections
  - ...
- New approaches
  - Parallelization-friendly languages (Haskell, Erlang, ...)

How

- Old technologies
  - Multi-processing
  - Multi-threading
- New technologies
  - OpenMP
  - Threading Building Blocks
  - Concurrent Collections
  - ...
- New approaches
  - Parallelization-friendly languages (Haskell, Erlang, ...)

Outline

Gaudi

LHCb: Gaudi.Parallel

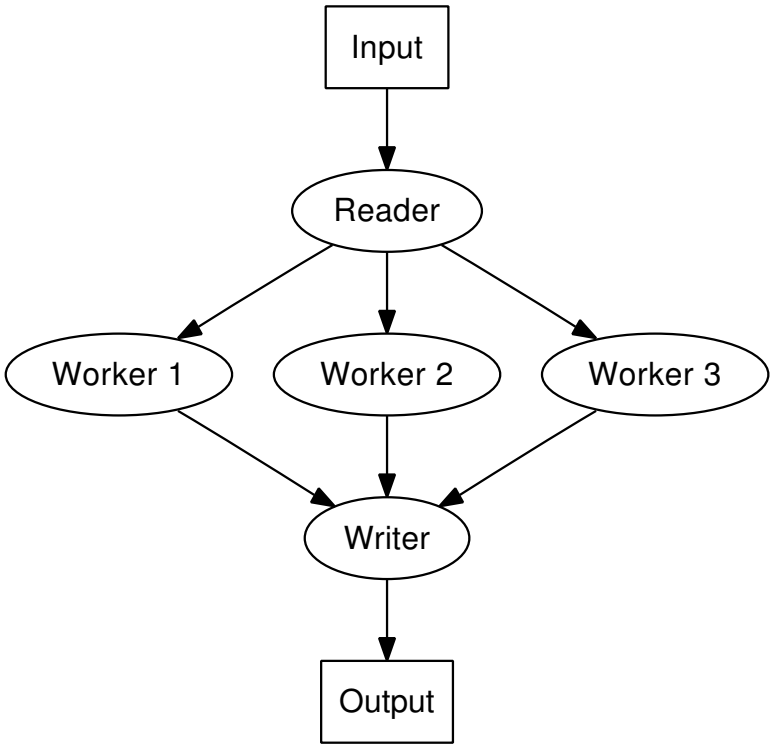

Atlas: AthenaMP

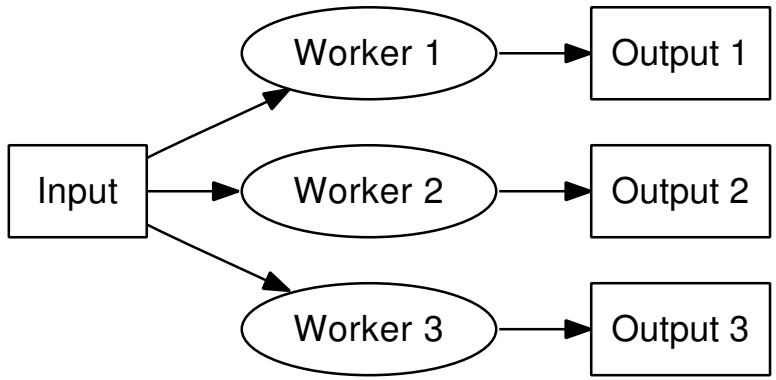

Plans: merge the two approaches taking the advantages of both

Gaudi

LHCb: Gaudi.Parallel

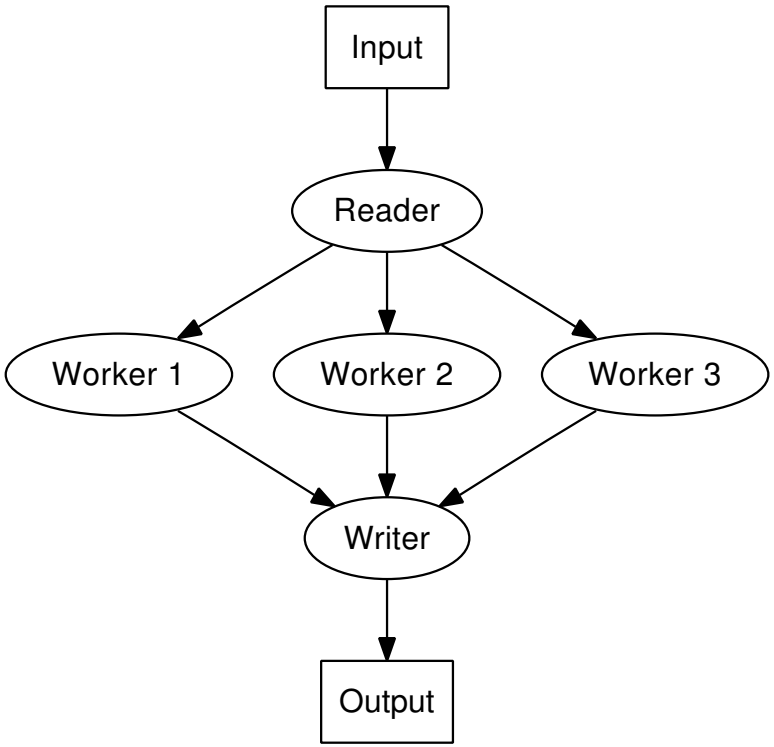

Atlas: AthenaMP

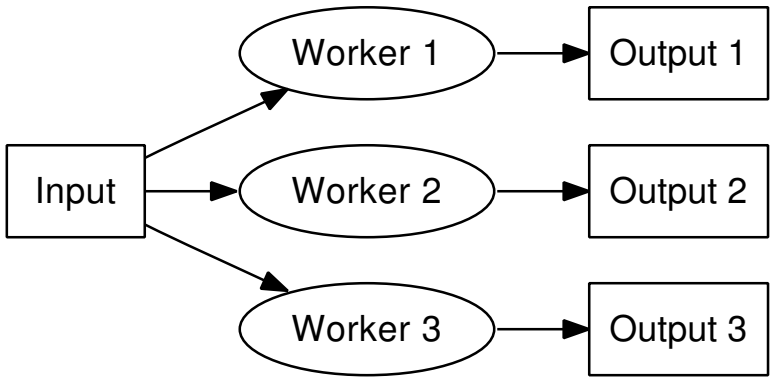

Plans: merge the two approaches taking the advantages of both

In-core Optimization

- Tools available that integrate in Gaudi (D. Kruse, K. Kruzelecki)
- We need guidelines to help developers to
  - understand the data
  - validate the changes

Outline

How can we use it?

- Multi-Core Laptops and Desktops
  - Can already use the parallel Gaudi and Athena
- Grid
  - Not ready: the batch system are tuned for the old way
  - New policies must come, we must be ready

How can we use it?

- Multi-Core Laptops and Desktops
  - Can already use the parallel Gaudi and Athena
- Grid
  - Not ready: the batch system are tuned for the old way
  - New policies must come, we must be ready

Summary

- Virtualization
  - A solid infrastructure is required to be usable
  - The cost in terms of performances must be taken into account
- Multi-Core
  - Getting ready to use multi-core computing elements
  - Support needed from the batch systems

Thanks to the people who helped me to prepare the presentation: E. Bonaccorsi, C. Jones, A. Papanestis, T. Ruf

Summary

- Virtualization
  - A solid infrastructure is required to be usable
  - The cost in terms of performances must be taken into account
- Multi-Core
  - Getting ready to use multi-core computing elements
  - Support needed from the batch systems

Thanks to the people who helped me to prepare the presentation: E. Bonaccorsi, C. Jones, A. Papanestis, T. Ruf

Summary

- Virtualization
  - A solid infrastructure is required to be usable
  - The cost in terms of performances must be taken into account
- Multi-Core
  - Getting ready to use multi-core computing elements
  - Support needed from the batch systems
- Can we use both or do we have to make a choice?

Thanks to the people who helped me to prepare the presentation: E. Bonaccorsi, C. Jones, A. Papanestis, T. Ruf