Coherence JMX Management Framework and Reporter

London Coherence SIG

January 27, 2011

Everett Williams

Senior Technology Director

SL Corporation

Everett.Williams@sl.com

Agenda

- Introduction/Disclaimer
- Acquiring Coherence Management Data
- Coherence JMX Reporter
- Next generation Coherence Management
- Extending Coherence with custom MBeans.

Privileged and Confidential

Acquiring Coherence Management Data

- Issues and Concerns
- Original Design and Optimization Patterns
- Coherence 3.4 upgrades
- “Real World” improvements

Privileged and Confidential

Issues and Concerns

- Data Accuracy
  - Time skew of data – The number of milliseconds between the first data element and the last data element of a snapshot.
  - Data Latency – How current is the data.
  - Sampling Frequency
- Capacity Consumption
  - MBean explosion
    - Number Caches * Number of Nodes (n-squared).
  - Peak vs Average consumption
  - Sampling Frequency

Questions?

Acquiring Management Data (original)

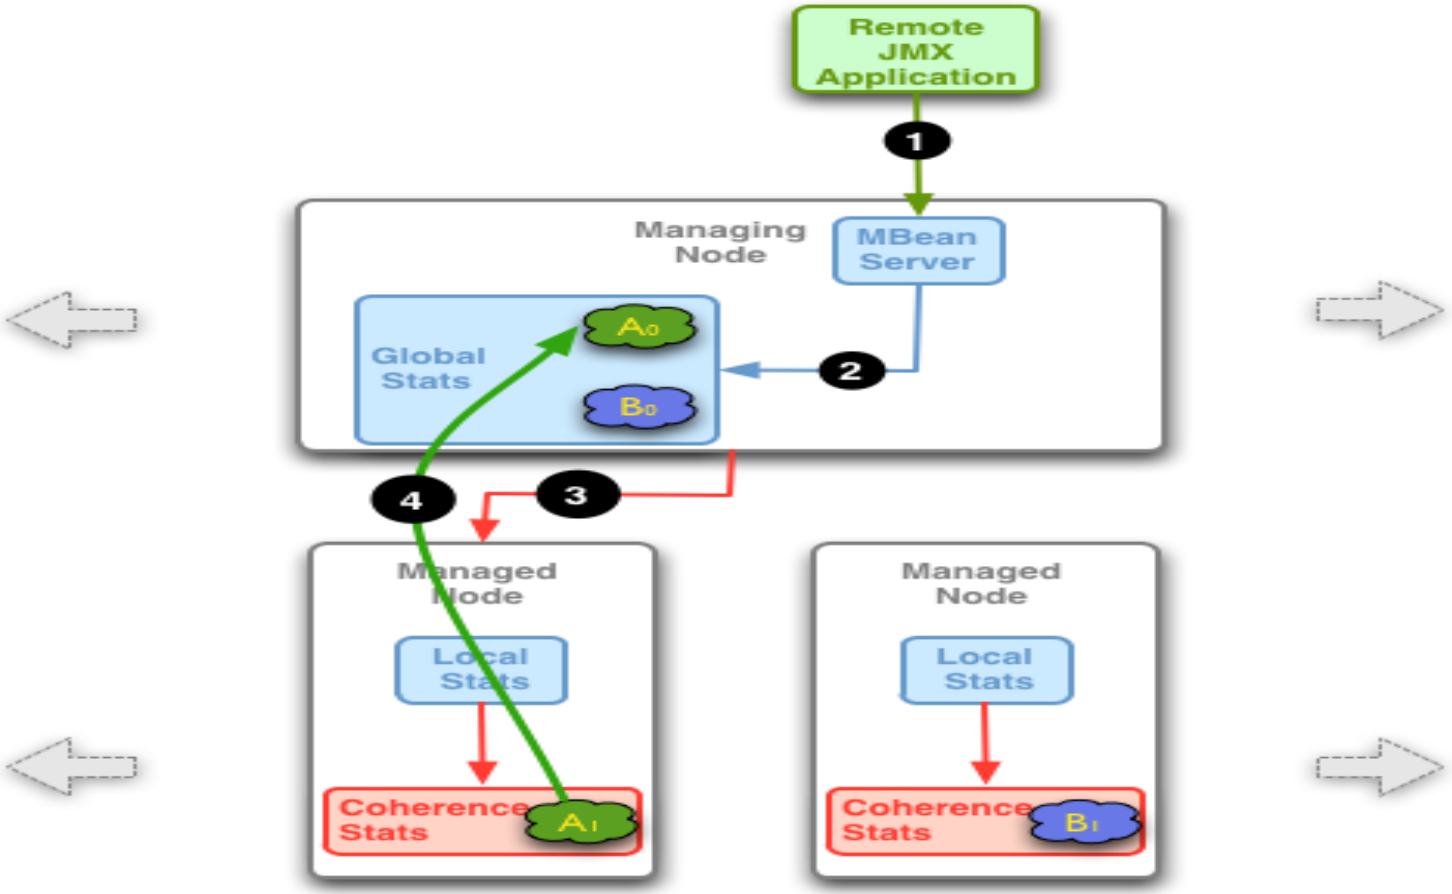

Privileged and Confidential

Issues with Original (refresh on demand)

- Current Data
- High Time skew of data
- N-squared communications
- Long time between sampling frequencies.
- Longer time to incident notification.

Privileged and Confidential

Acquiring Management Data (pre-fetch)

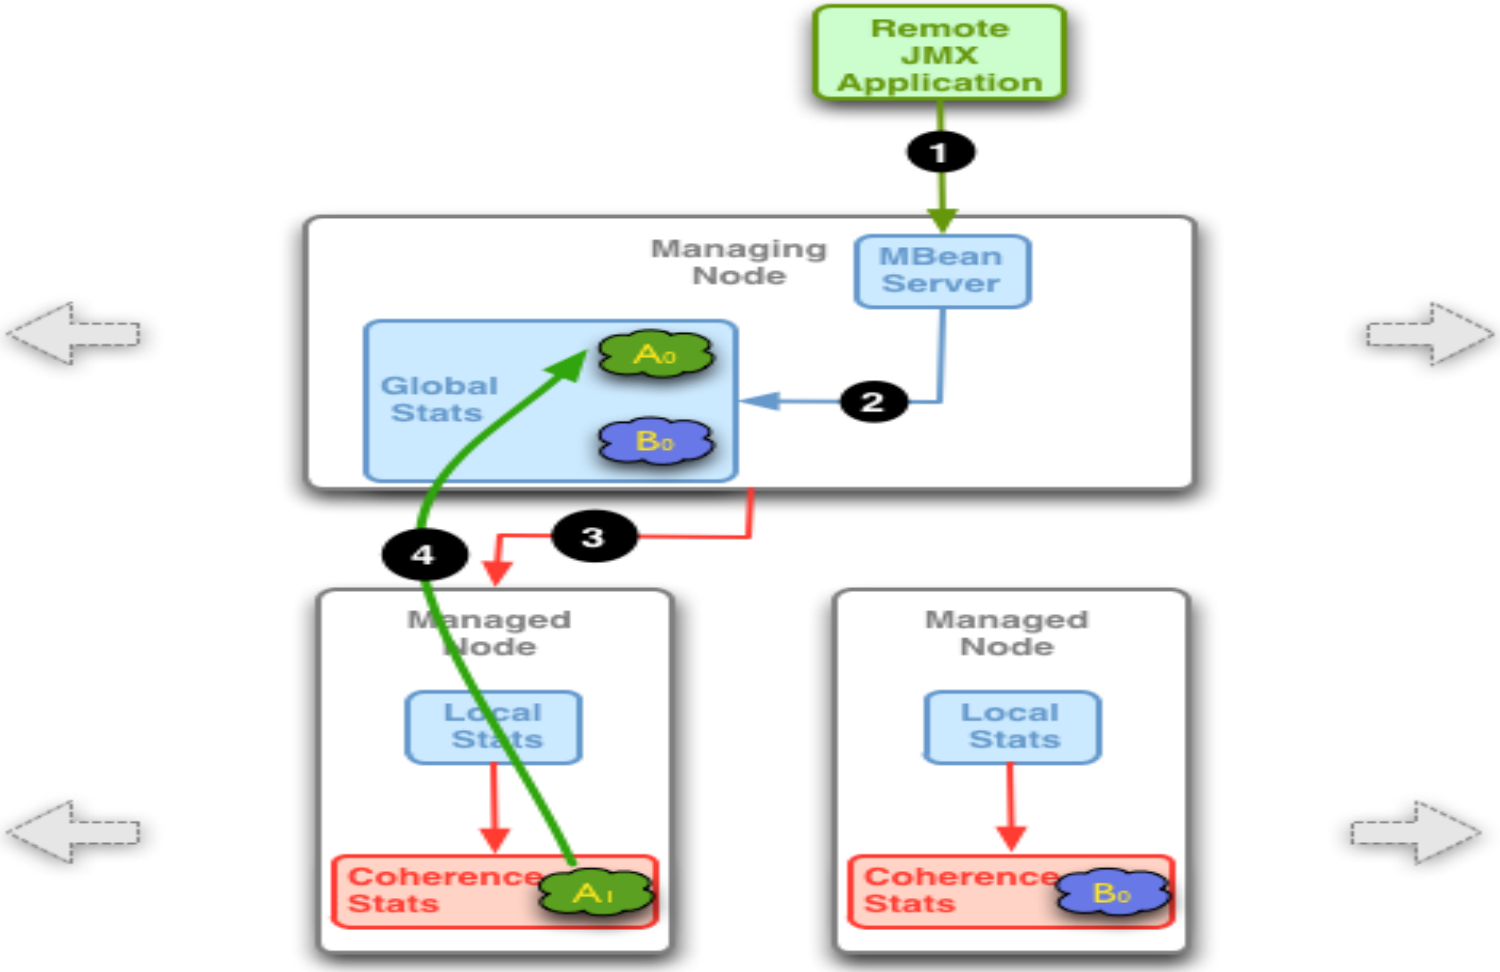

Privileged and Confidential

Acquiring Management Data (pre-fetch)

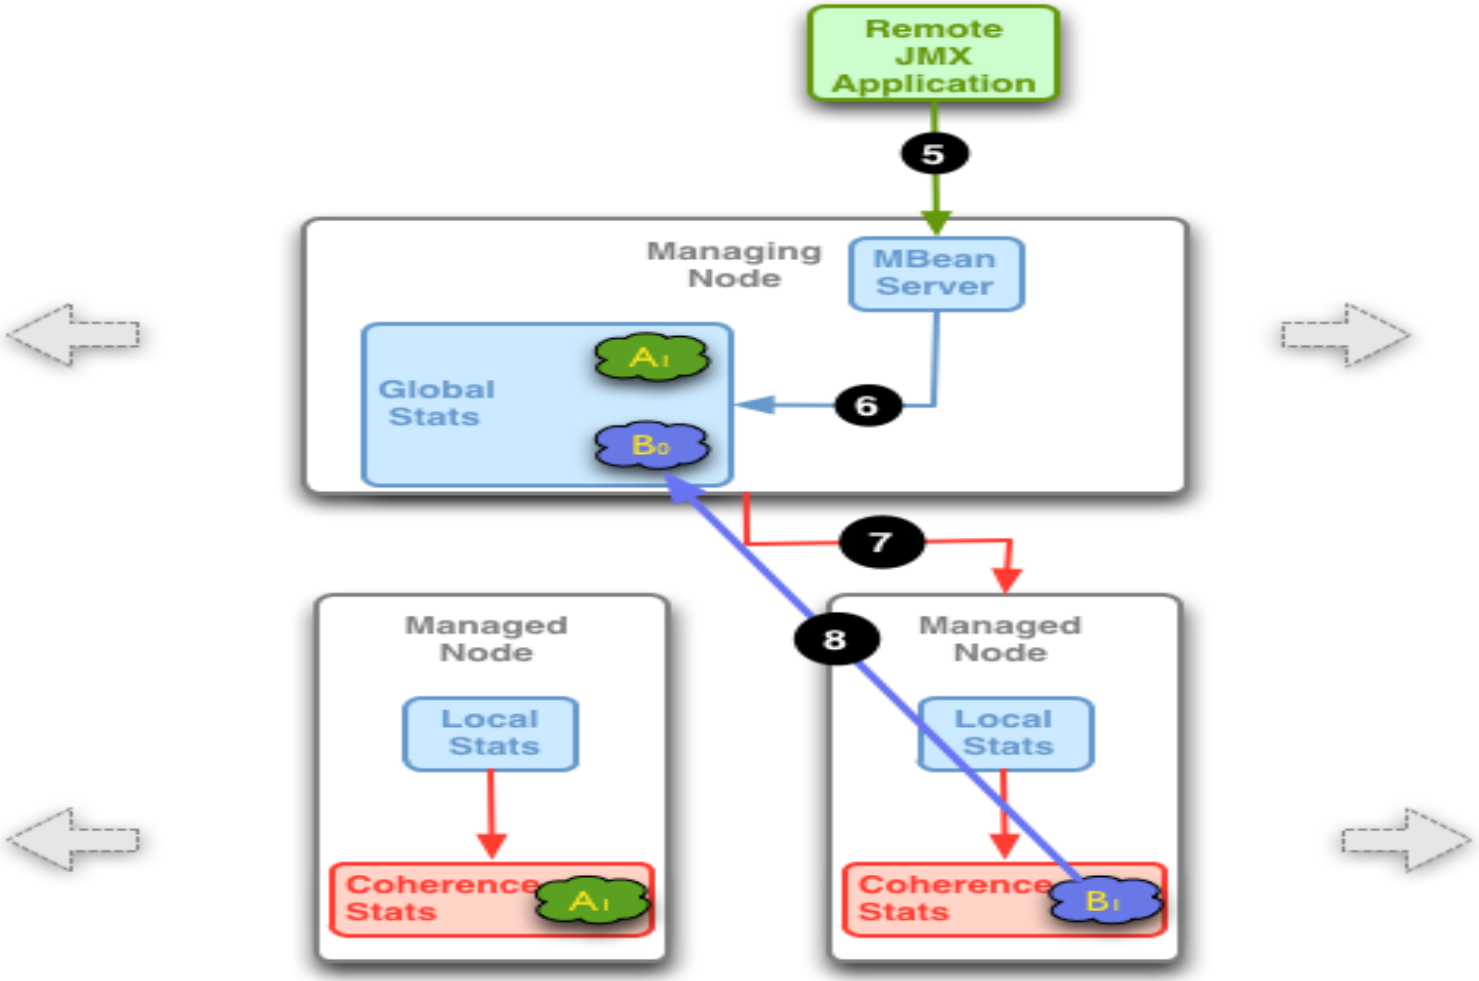

Privileged and Confidential

Acquiring Management Data (pre-fetch)

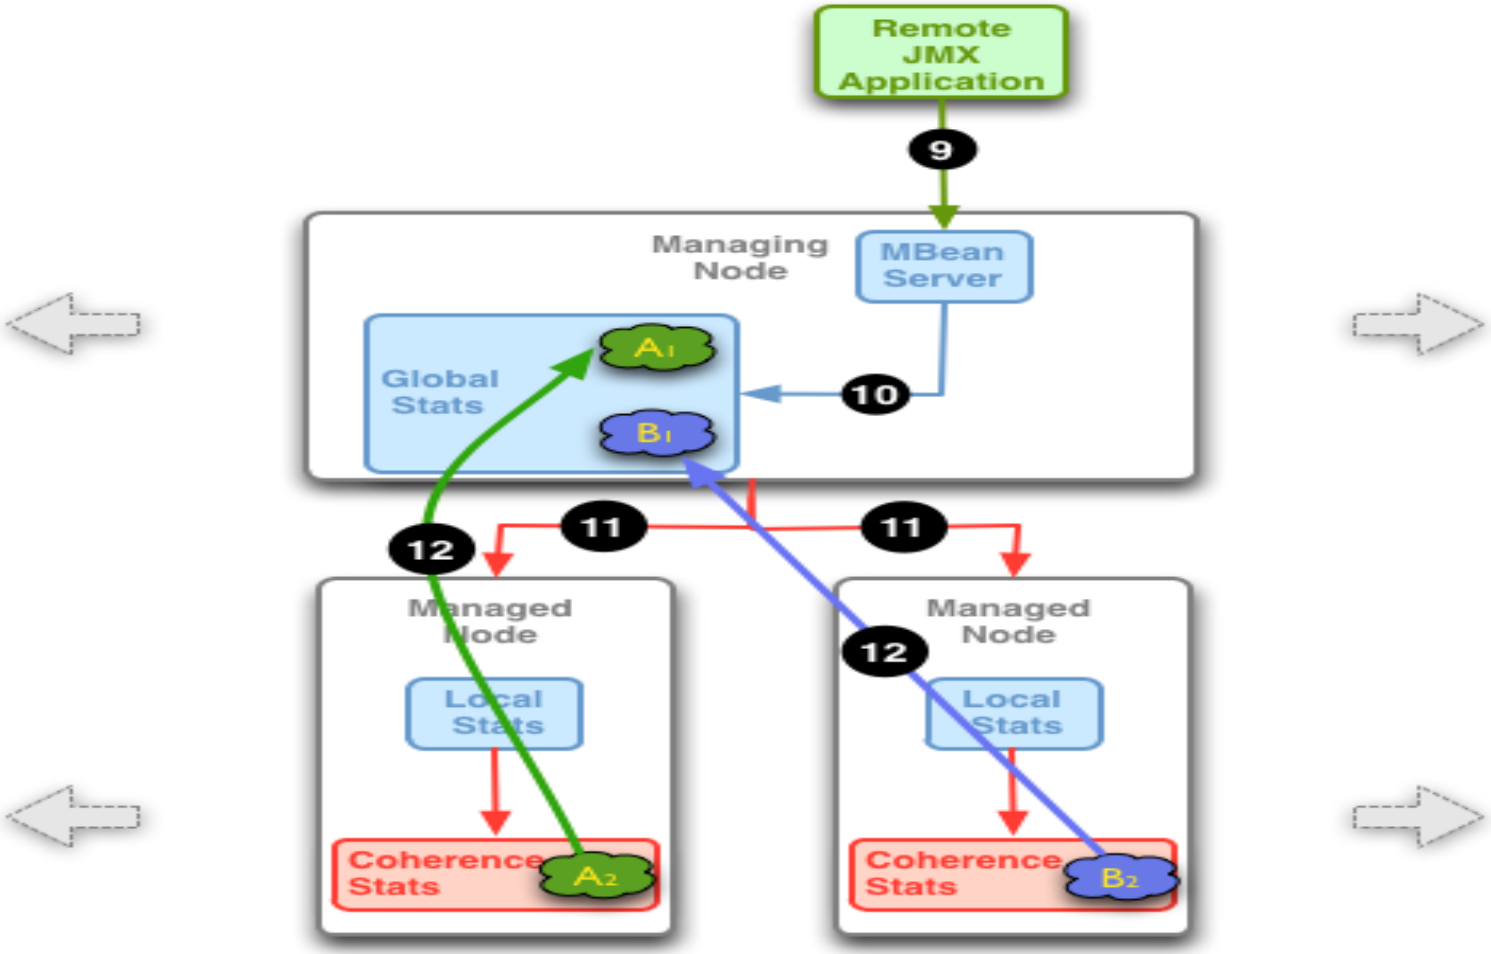

Issues with prefetch

- It takes at least one cycle to determine the pattern
- Multiple usage patterns on the same node significantly impacted performance
- Increased the maximum load on the coherence nodes (spike)
- Increased the load on the managing node (client).
- n-squared communications.

Acquiring Management Data (real world)

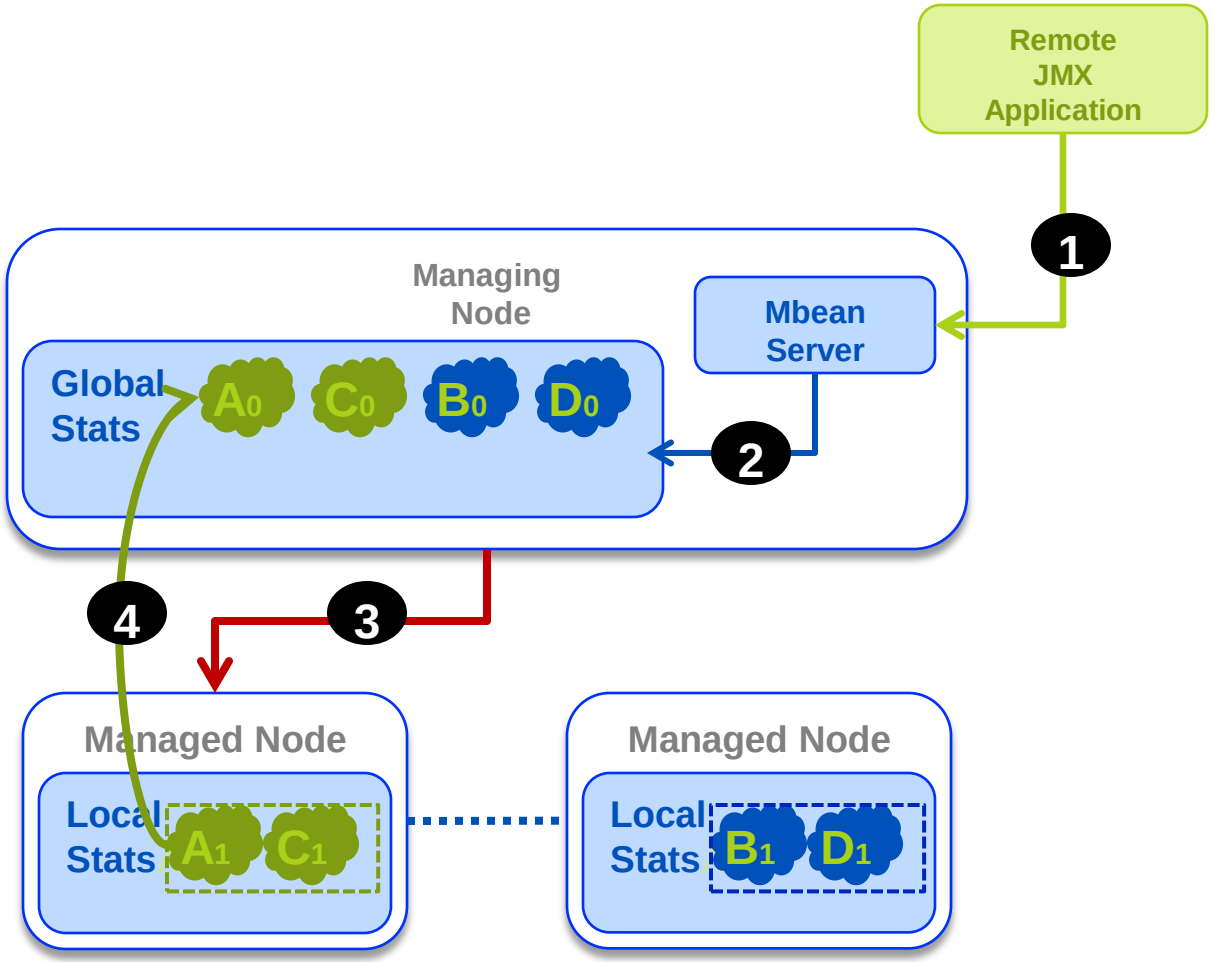

Privileged and Confidential

Acquiring Management Data (batch)

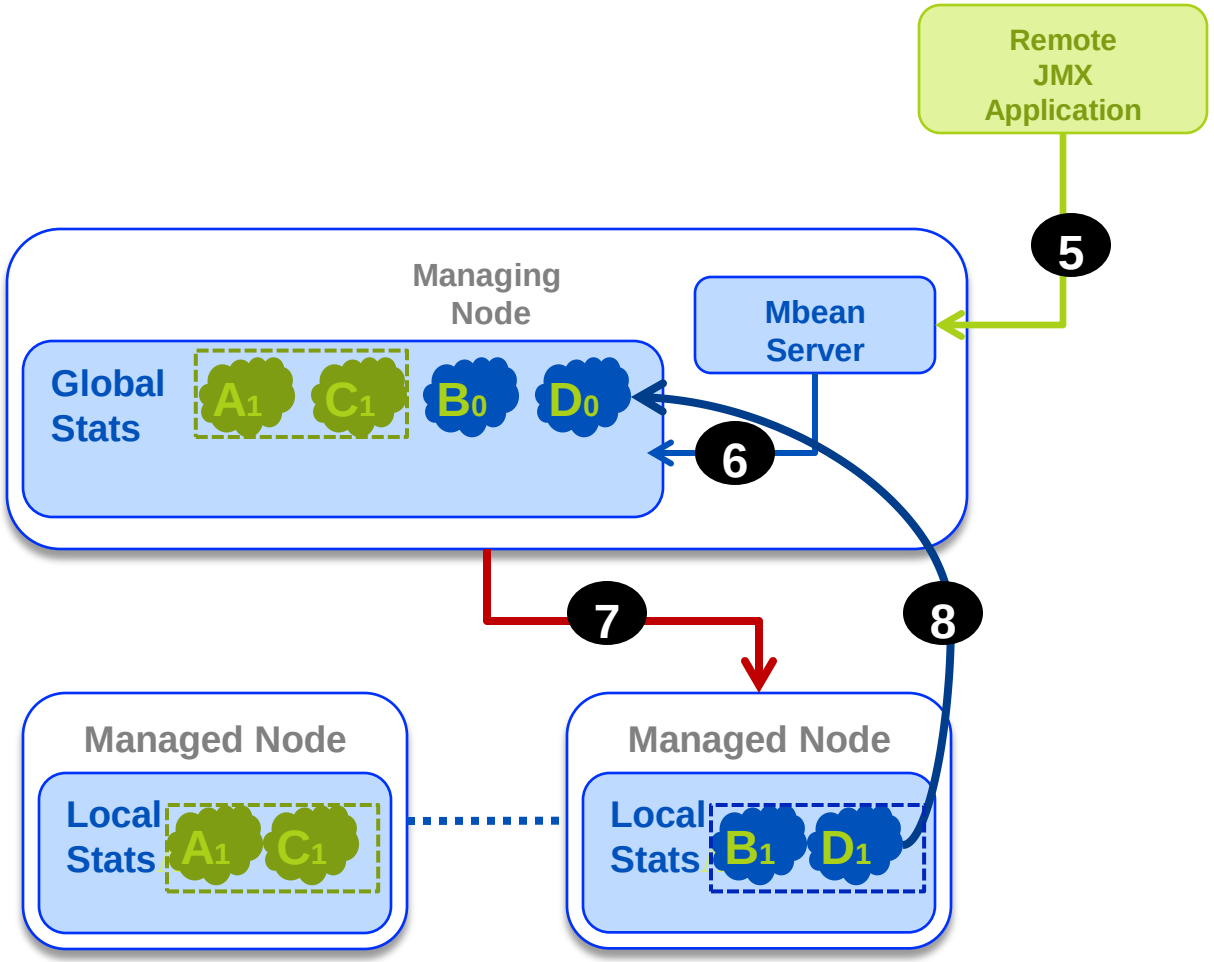

Privileged and Confidential

Acquiring Management Data (batch)

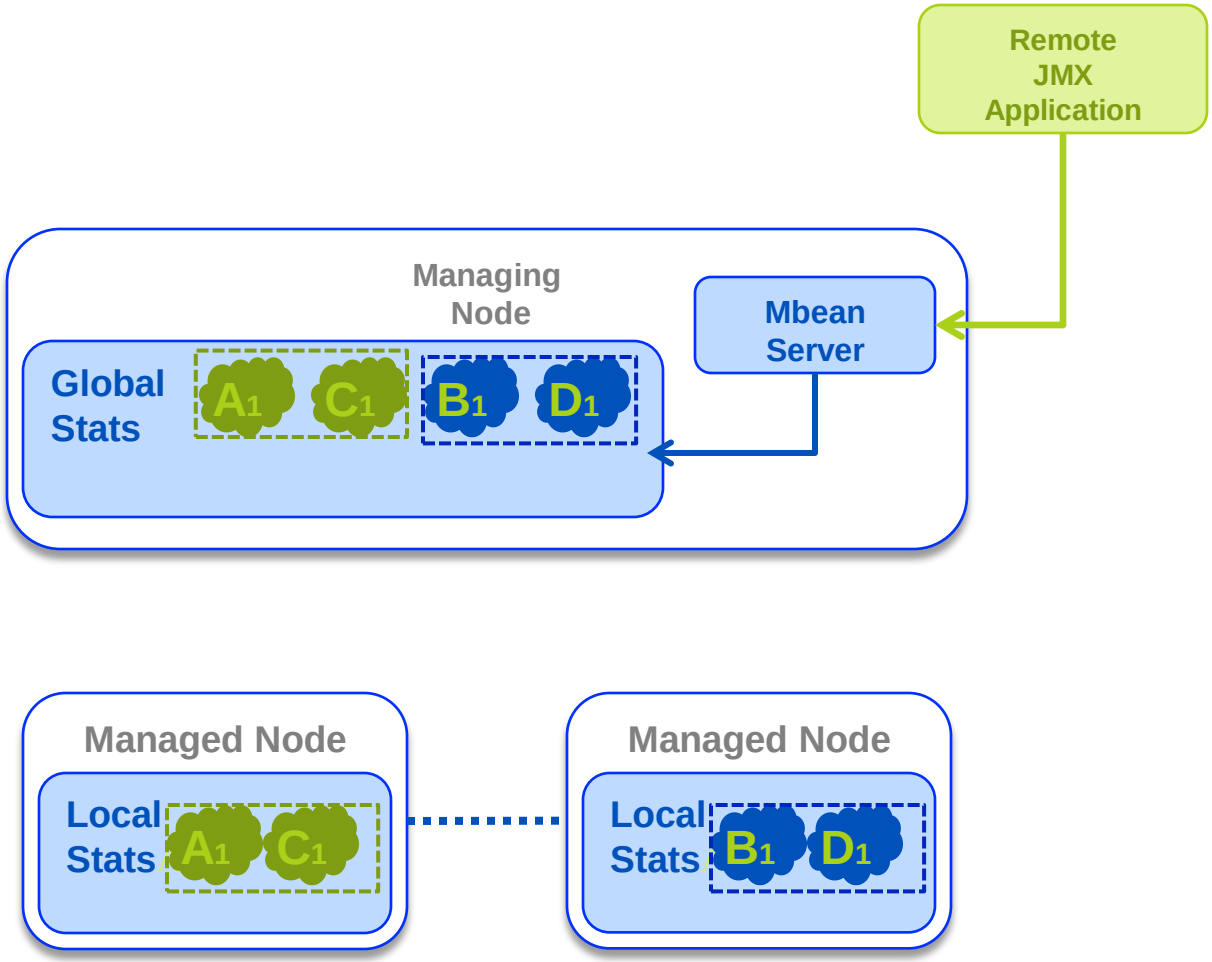

Privileged and Confidential

Batch Retrieval

- Optimized for snapshots
- Minimizes Time Skew of data from each node and overall cluster.
- Independent of refresh policy
- Order N (node count) communications per snapshot.
- Significantly reduces the load on the cluster. ~(factor of 3)
- Significantly reduces the time from incident to notification.

Questions?

JMX Reporter

- Original design and goals
- Configuration
- Starting, Stopping, Status
- Data Generated
- Reporter Features

Original Design and Capabilities

- Designed for POC teams.
- Zero additional infrastructure
- JMX Capture Utility
- Central and Distributed execution model.

Configuration (Report Definition)

- JMX Query
- Simple Join
- Row operations (add, subtract, divide)
- Aggregations
- Delta Operations
- Row Filters

Configuration (Report Group)

- Report Groups are a set of report definitions that are run together.
- There is one report-group per node.
- Out of the box report groups
  - Report-all.xml – all the reports will be executed
  - Report-group – standard reports.
  - Report-web-group.xml - Standard Reports + coherence web reports.

Standard Report Group

- Cache Usage (Hits, gets, puts, etc)
- Service Information (Tasks, Requests, etc)
- Network Health (Publisher/Receiver Success Rates, packet information)
- Network Health Detail (by node)
- Node listing
- Management Performance.
- Proxy Node and Proxy Connection.
- Each Row of Data is stamped with a batch number, a report time, and refresh time.

Starting and Stopping

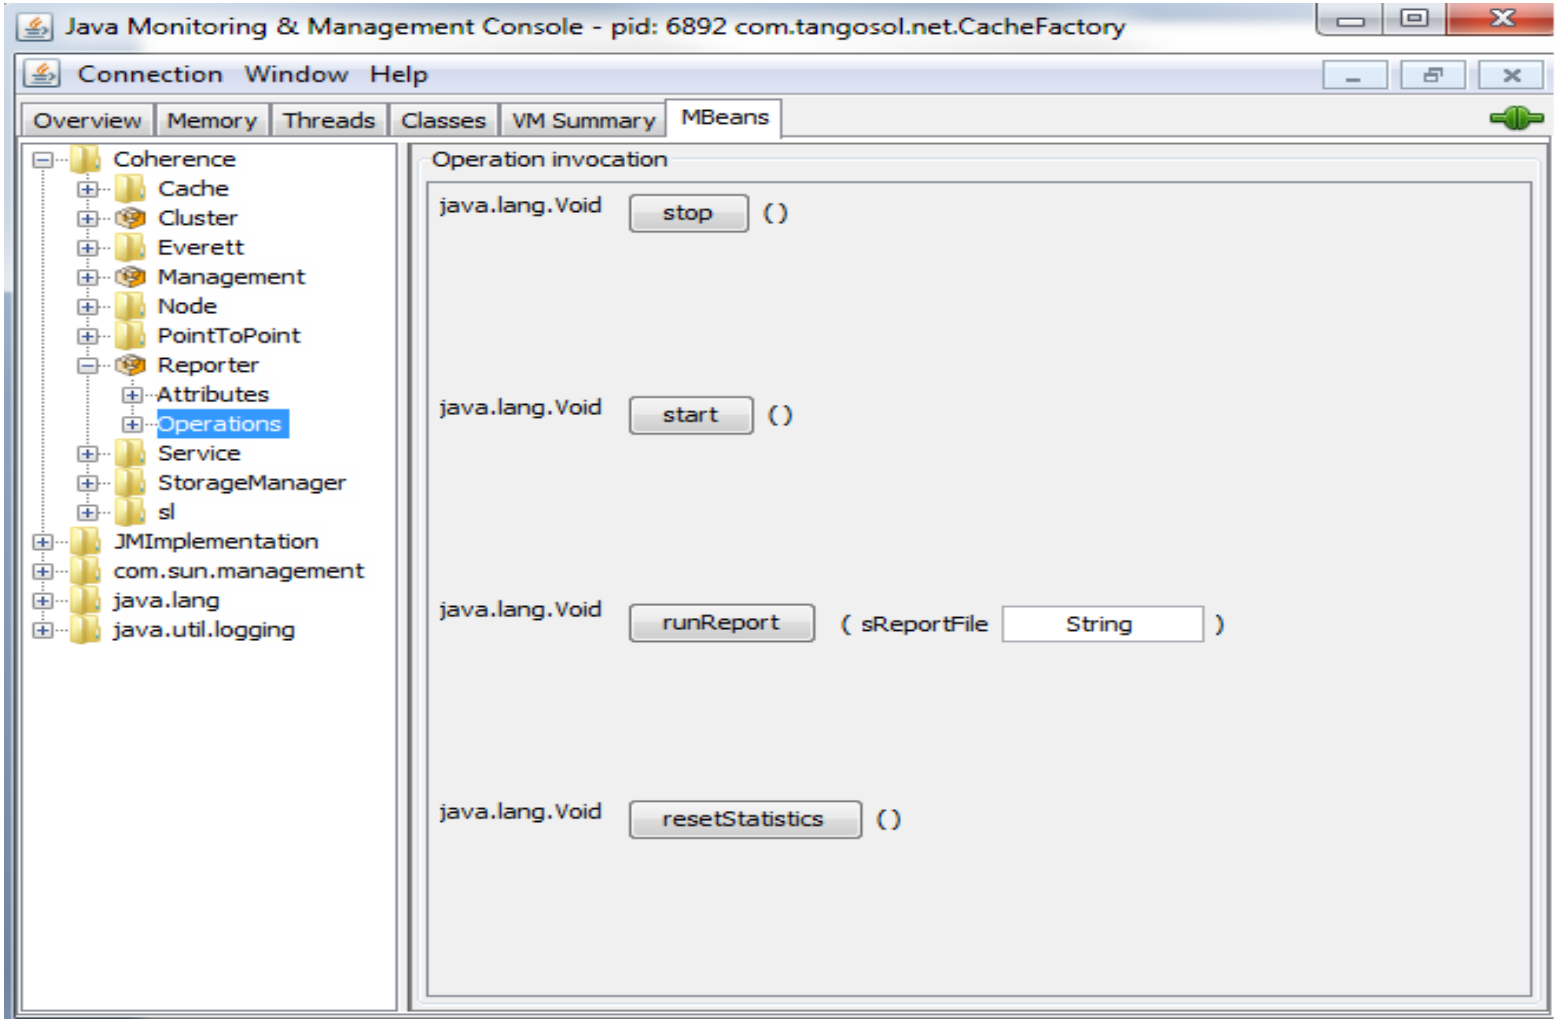

Command Line

- AutoStart
  - -Dtangosol.coherence.management.report.autostart=true
- Report Group
  - -Dtangosol.coherence.management.report.configuration=<file>

Reporter Status

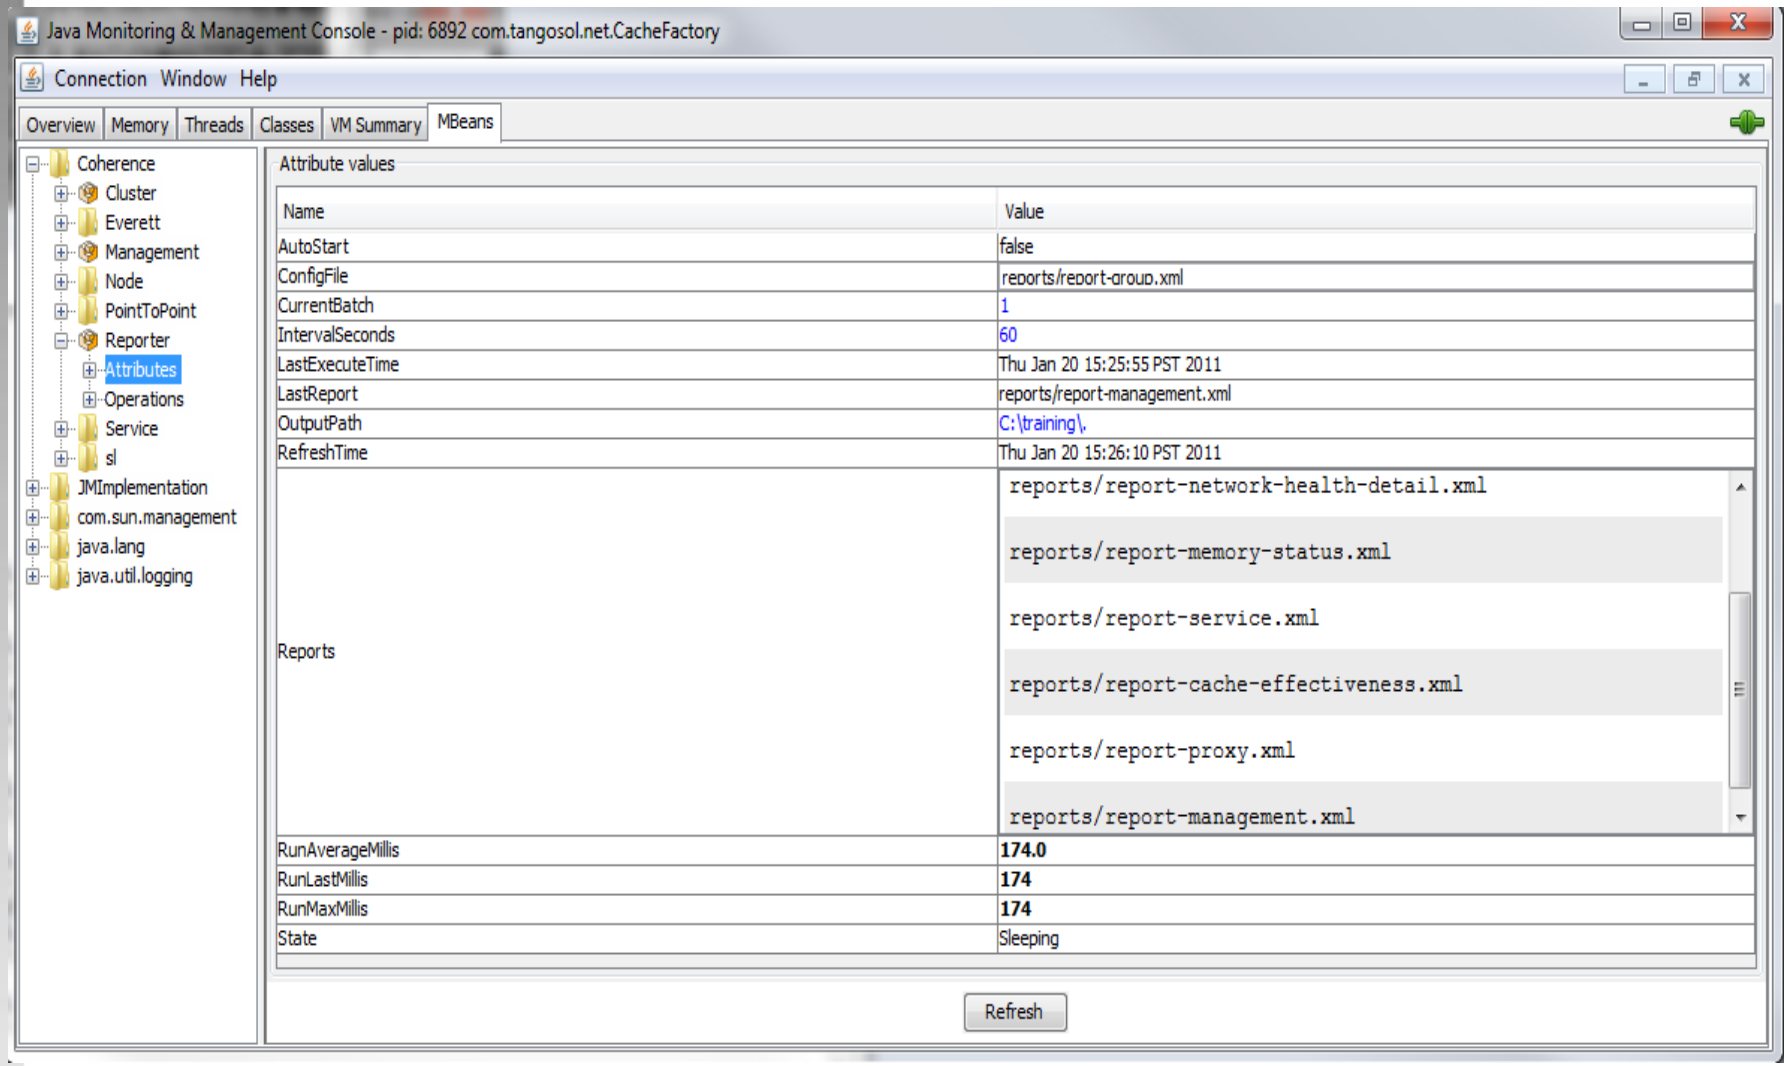

Reporter Status

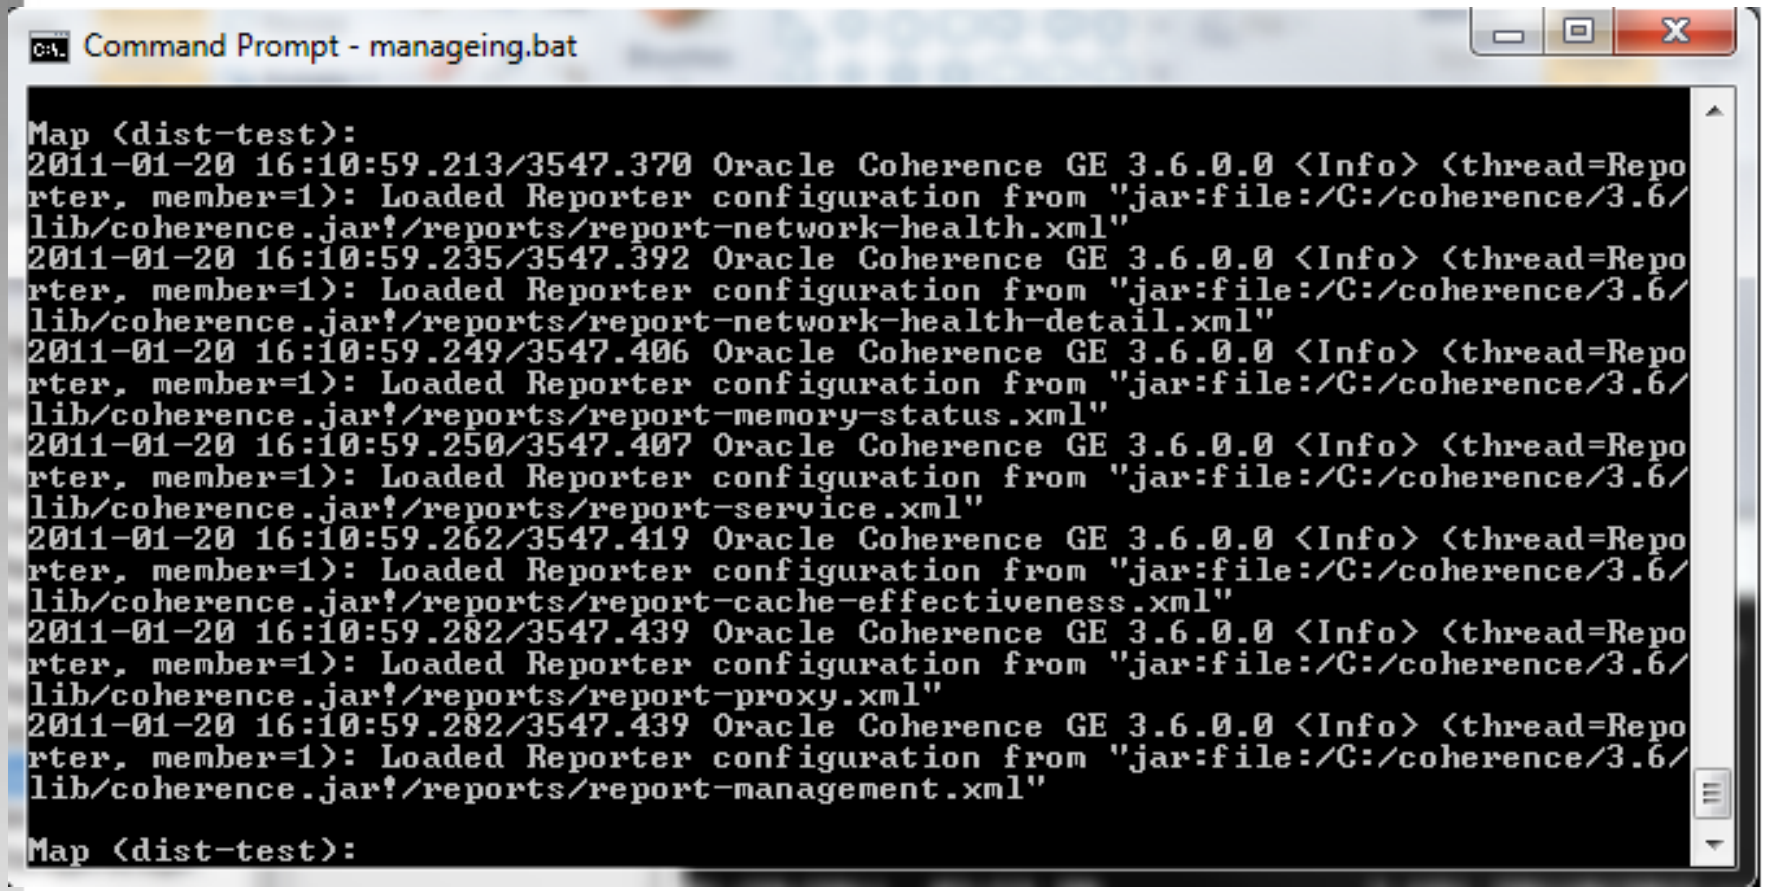

Reporter File Listing

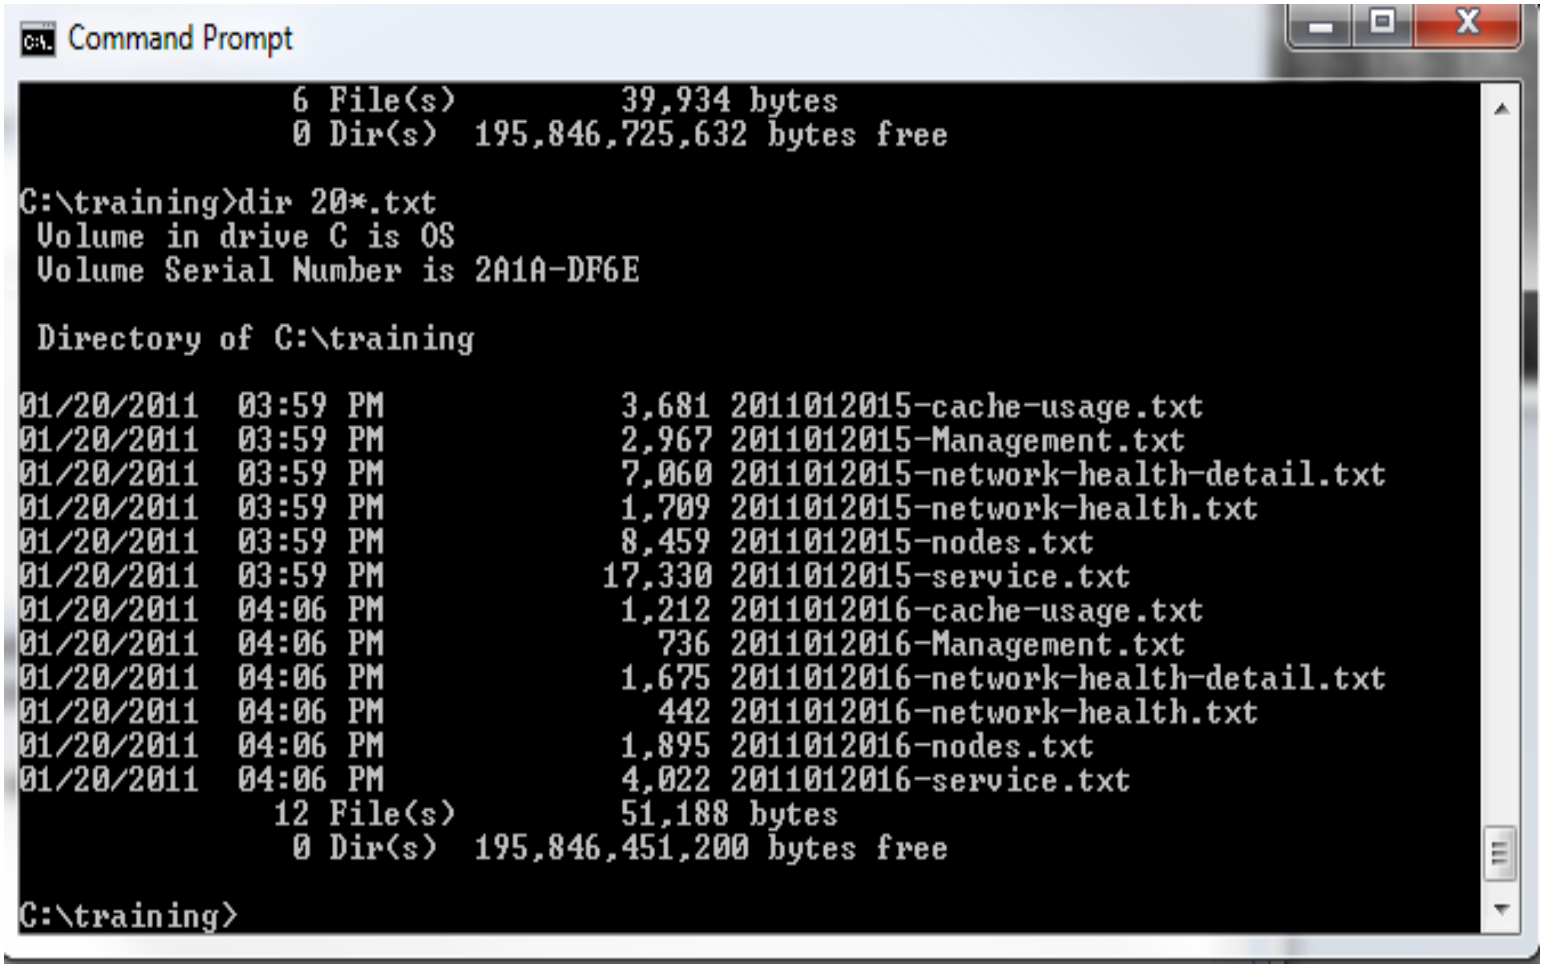

Reporter Links

- Analyzing Reporter Data
  - http://download.oracle.com/docs/cd/E15357_01/coh.360/e15723/manage_reportercont.htm#CIHIHEHG
- Creating a Custom Report
  - http://download.oracle.com/docs/cd/E15357_01/coh.360/e15723/manage_createcustreport.htm#CIHDCCFE

Coherence JMX Reporter

- Proprietary XML
- No History
- Simple Data Management
- No Analytical Capability – No trending, history, etc
- Single Threaded, Single Timer

Questions?

Notifications

- Were added in Coherence 3.6 as support for custom platform Mbeans
- Support Notifications generated by Standard, Dynamic or 3rd party (replicated) Mbeans
- Currently only two Coherence Specific Notifications (member left, member joined.

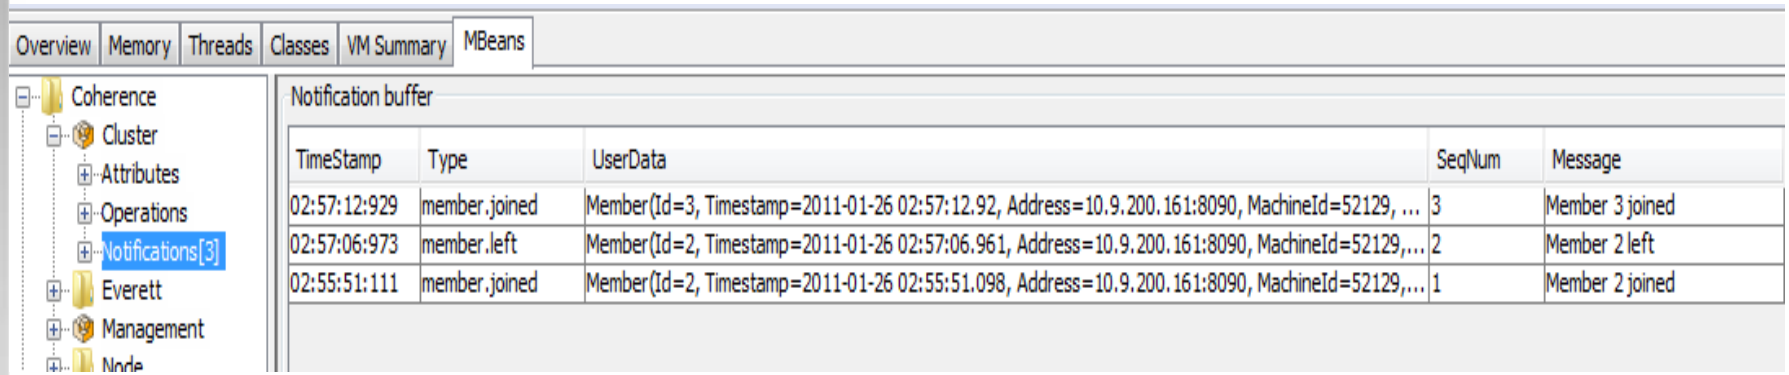

Privileged and Confidential

Notifications Impact

- Usage Patterns
  - Point in time metric monitoring and thresh alerting
  - Event Notification
- Significantly increase the scalability of Coherence Monitoring
- Does not eliminate the need for Sampling (polling)
- Should be used with caution.
  - Over Notification can be worse than polling
- Algorithms to process and interpret the notification streams.
- Reduce the time from the incident to when the support team is notified.

Questions?

Extending Coherence with Custom MBeans

- Example Usage Patterns
  - Application Performance metrics
  - Centralizing 3rd party MBeans
  - Extending “missing” functionality
- Integrating Custom MBeans
  - Standard MBeans and Notifications are simple to implement
  - Integrating with coherence is as simple as modifying an XML file.
  - Caution should be used to avoid impacting performance with too many measures.

Extending Coherence with Custom MBeans

- Benefits of Custom MBeans
  - Provides a consolidated view of the impact of application usage on coherence resources.
  - Consolidating infrastructure component (container, etc) statistics helps spot systemic bottlenecks
  - Leverage 3rd party analytic tools for visualization and management of the system.

Questions?