

Practical Considerations When Instrumenting Applications With JMX

Tom Lubinski
SL Corporation
Corte Madera, CA
March 18, 2008

Abstract – There are many sources of information about best practices using JMX, focused primarily on common usage conventions and design patterns. However, the use of JMX in custom application monitoring systems brings with it an array of unique challenges involving the aggregation, analysis, and visualization of real-time data. This paper discusses practical issues regarding data content and design of MBeans for optimal use in such systems. Some commonly overlooked requirements are explored and recommendations made for minimizing client-side development efforts.

I. INTRODUCTION

With any new technology, best practice documents are invaluable in helping developers avoid common errors and to design quality systems. Much literature is already available regarding best practices for using Java Management Extensions (JMX) in monitoring and management applications. Popular J2EE application servers, such as BEA WebLogic and JBoss, have used JMX for years to manage and monitor the health and status of their many components.

These large scale systems were built using an early version of Java (1.4) and add-in libraries of JMX classes. The extra steps involved in using JMX limited its use to systems in which the benefits of exposing monitoring and management information outweighed the cost of developing and supporting the additional code – JMX was simply not in common use.

With the release of Java 1.5, JMX is built-in and readily usable in even the smallest of applications. It is now a simple task to instrument nearly any application and expose important monitoring metrics. Custom applications, involving many processes and multiple middleware components, may be effectively managed from a remote console. As a result, there has been an explosion of data exposed through JMX and available for analysis and presentation.

As use of JMX expands, one would expect that there would be mistakes made and lessons learned the hard way. Software developers are usually much more familiar with their own application domain and are not often experts in monitoring and management tools.

The author of this article and SL Corporation have over 20 years of experience with monitoring and visualization applications, with particular expertise in Java. The company's Enterprise RTView product has been specially adapted to deal

with real time data produced by JMX-enabled applications and has features to compensate for many overlooked requirements.

This article discusses errors that have been seen repeatedly in JMX implementations regarding the content and design of data structures known as MBeans. Common JMX best-practice knowledge is briefly reviewed to provide some initial context. This is followed by a detailed discussion about issues that arise when custom application monitoring requirements grow to include aggregation, analysis, and visualization in real time. Recommendations are offered that that may help users make optimal use of JMX in these situations.

II. COMMON BEST PRACTICES

JMX is a very general solution framework and does not define specific monitoring or management data structures. This puts the burden on developers to establish conventions themselves in order to consistently extract and process information so it can be analyzed and visualized.

Best practice documents for JMX typically suggest adherence to common standards regarding naming conventions, data types, deployment constraints, portability, and so on. These suggestions fall into several categories:

A. MBean usage conventions

Specific conventions should be applied to MBean names and JMX data types in order to implement MBeans that are portable, i.e., commonly available JMX client applications can handle them without issue. Generally, this means one should follow the guidelines for using standard domain and key names (such as `type=` and `name=`) and `OpenMBean` data types instead of custom Java classes (for ease of deployment). These suggestions are presented in depth in SUN's discussion of JMX best practices[1].

MBean names should be predictable and used in a consistent manner when representing a hierarchy. Most importantly, one should define the same attribute schema for all beans of the same type or at the same level in a hierarchy. It is obvious that a violation of this principle will result in complications for client applications.

Use Standard MBeans wherever possible; these are the simplest to implement and the easiest to maintain, since JMX implicitly understands data types and runtime behavior simply from the Java source directives. There are situations where Dynamic and/or Model MBeans may be necessary, but their use should be kept to a minimum [1].

B. Follow established design patterns

Use well-established design patterns for defining MBeans. Several of these are described in *Design Patterns for JMX and Application Manageability*[2]. Often, the motivation behind these patterns is to shield users of JMX MBeans from the details of their implementation or even their attribute structure. This permits the developer to change the details of the MBean without affecting client applications.

Other patterns are useful in maintaining separation between the monitoring and management code contained in the JMX MBean, and the business logic in the application. There are significant advantages in terms of maintainability if the application being monitored has little or no knowledge of how the monitoring is done.

There is also much written about Aggregator beans as a way to minimize the number of MBeans that a client has to connect to and query. When the number of beans reaches into the thousands, performance issues come into play. The Aggregator pattern is one way to improve performance by minimizing MBean access.

C. Model after established systems

The most popular J2EE application servers make extensive use of JMX as a tool for managing their internal functions. These very large systems provide solid ground for testing implementation techniques for JMX MBeans. As such, they provide excellent examples to follow.

In the administration system for the BEA WebLogic Server, there can be over 1000 different MBean instances. MBeans have been developed to address just about every variation of monitoring and management problem. By exploring the techniques used in this system, much can be learned and applied to your own requirements. The BEA JMX system is especially complete when it comes to the use of notifications for monitoring system performance [3].

The J2EE Management Specification JSR-77 [4] defines a Management Model and provides useful guidelines for data types and implementation patterns. Effective management and monitoring involves states, statistics, metrics, relationships, and more. Understanding these information structures can be helpful in developing quality systems.

III. MONITORING, ANALYSIS, AND VISUALIZATION ISSUES

The best practices outlined above are useful, but they don't fully address issues that come up in large scale monitoring applications involving visualization and, in particular, dynamic analysis of real-time data. Requirements, such as trending or slice-and-dice data analysis, are often afterthoughts. It is important to understand these issues up-front and develop an implementation plan that minimizes the work required to present the data. Often, this can influence the design of the MBeans developed for such an application.

The best way to illustrate this is with a concrete example. In this case, consider developing a monitoring system for a

sample application, a simple message switching system. This example is representative of many applications. Similar data structures are relevant for monitoring performance of routers, message boxes, caching systems, object databases, even CEP engines. Shared among them is the desire to use JMX for instrumentation.

Our sample device supports multiple channels for message traffic and collects metrics and statistics about its operation in real time, so performance can be optimized. Here is a simple data structure containing information that might be exposed via a JMX MBean for a single channel:

ChannelInfo Data

Channel	Msgs Sent	Msgs Rcvd	Bytes Queued	Max Bytes	Statistics
12	12549	9613	55040	2349128	Proc=124ms, Wait=37ms

The Total Msgs Sent/Rcvd counts would typically be represented as a long integer since the number could grow large over time. The memory usage could be stored as an integer since the Java memory limit on many machines is 1 GB, so 32 bits will handle that effectively. The other data – Channel ID and Statistics – are stored as strings.

This MBean would be assigned a domain name, like “ChannelManager”, and a key, like “type=ChannelInfo”. Each channel would create a unique MBean with two additional key components, “server=XX” and “channel=NN”, to indicate the server on which the channel is running and the ID of the channel itself. The full name of the bean used to return information about channel 12 on Server1 would be:

ChannelManager:type=ChannelInfo,server=Server1,channel=12

Additionally, we might have other information related to configuration, perhaps not as dynamic. This might include an IP Address and Port, and a count of active connections:

NetworkInfo Data

Channel	IP Address	Port	Connections
12	192.168.2.103	4901	17

From a developer perspective, the information collected here is obtained from a different source and, as such, it is natural to have a separate MBean. It would have a different type key, “type=NetworkInfo”, indicating it is network-related data. The full name for this bean might be:

ChannelManager:type=NetworkInfo,server=Server1,channel=12

Thus we have two simple MBeans, a ChannelInfo and a NetworkInfo bean. These map closely to the internal data structures of the system, so it is easy and quick for the developers to implement.

At this point, things are looking good. Lots of information is available for monitoring. Typically, an HTML page is provided as part of the system to view each of the MBeans. The system runs, the data shows up in the HTML page and everyone (especially marketing) is excited about the new

application monitoring capability available with this new version of the company's product.

A. *Not so fast ... there clouds on the horizon*

The first indication of a problem shows up when customers attempt to monitor the system in a useful way against live data. At first, looking at each channel via the HTML interface is exciting... there is so much data to explore and users see things they've never seen before about the behavior of the system.

But looking at one bean at a time in an HTML page gets old really fast, especially once you have more than just a few channels. Imagine how difficult this is if one of the servers has 100 channels running on it. It is practically impossible to extract useful information from the system this way.

It quickly becomes apparent that to fully understand the workings of the system, one must be able to perform calculations on the incoming data in aggregate. There must be a way to sum messages counts across all channels, or take an average across servers. Metrics can be dumped to a database for later analysis, but this is not a real-time solution by any means.

One solution is to take advantage of the advanced capability provided by JMX to access multiple MBeans using wildcard "*" syntax. In other words, make a request for the names of all beans matching a certain pattern, such as:

```
ChannelManager:type=ChannelInfo,server=Server1,*
```

This request will return the names of all ChannelInfo beans on Server1 for all channels. The data from each can be compiled into a single table, one row per bean. This seems like it will work, but there is a big problem lurking.

B. *Identify yourself...or else*

In our sample MBean, it seemed natural for the developer to include the Channel ID in the data – on a single server, the internal data structures contain the ID of the channel being implemented. However, since all channels on a server live within the context of that server, it did not seem reasonable to include the name of the server as part of the data.

In fact, from the developer's perspective, it made sense to minimize the data to be transferred and leave out the server name – since the bean name itself contains the server info, it shouldn't be necessary to include it in the data.

However, look what happens when the data from each channel MBean on each server is gathered into a single table:

ChannelInfo Data – Multiple Channels

Channel	Msgs Sent	Msgs Rcvd	Bytes Queued	Max Bytes	Statistics
12	12549	9613	55040	2349128	Proc=124ms, Wait=37ms
13	9456	8719	
14	4456	3398	
12	22315	17617	
13	9981	8871	

Note that the table contains a row of data from each MBean on all servers. In this sample, the Channel ID field contains the number 13 twice. This is because the channel 13 exists on more than one server.

The fact that the data does not identify the name of the server from which it came is a big problem. The name of the MBean contains the name of the server, but use of the wildcard syntax to reference all MBeans does not automatically provide information about the source. Without information about the server, the table produced here is useless. It seems like a simple oversight, yet this is probably the most common problem encountered when visualization and analysis is attempted on data exposed via JMX (and many systems that pre-date JMX).

After seeing this situation come up time and again, some products like Enterprise RTView have evolved to provide automatic ways to supply this essential information to the presentation and analysis layer. This is done by parsing keys contained in the MBean name and creating new columns if they don't already exist. With the "Server" column added, the table in our example above becomes:

Enhanced ChannelInfo Data – Multiple Channels

Server	Chan	Msgs Sent	Msgs Rcvd	Bytes Que'd	Max Bytes	Stats
Server1	12	12549	9613	55040
Server1	13	9456	8719	
Server1	14	4456	3398	
Server2	12	22315	17617	
Server2	13	9981	8871	

In the example above, the "Channel" column already exists so it is not needed. The "Server" column is not available but can be added from the name of the MBean sourcing the data.

In general, the problem can be avoided by properly identifying the source of the data in the original data table. In many of these systems, there are dozens of columns of data. Saving one or two columns is not very effective when you consider what has to be done on the client side. There may be a question as to whether it is good modeling practice to include extra data in the MBean (it is duplicated in the name). However, if the goal is to minimize development effort on the client side, then the extra information is helpful. In an Aggregator MBean, it would be essential.

RECOMMENDATION: When using * to collect data from multiple beans of the same type or multiple connections, include columns as attributes that identify the source – otherwise you cannot tell which bean it came from.

C. Calculating rates... it's harder than it looks

Once the data is made available and properly identified, the problem then becomes presentation. One common requirement is to plot metrics like Msgs Sent and Rcvd in a trend chart.

Obviously, raw Total Msgs data should not be plotted, since it will be continually increasing. The delta from one event to the next must be computed and used to plot the trend.

This problem is not unique to JMX – it is found in many systems that gather data in real time. It is simple to count the messages, export that raw data, and let the monitoring client deal with the job of calculating deltas and rates.

However, at the client end it is not that simple, particularly when data are gathered from multiple sources. The sample tables below contain data at two instants in time for just two channels, 12 and 13, on Server 1.

Time Instant 1

Server	Channel	Msgs Sent	Msgs Rcvd	...
Server1	12	12549	10317	...
Server1	13	9456	8755	...

Time Instant 2

Server	Channel	Msgs Sent	Msgs Rcvd	...
Server1	12	12936	10566	...
Server1	13	9467	8789	...

As long as the number of channels is the same at each time instant, the problem is not so difficult. New values for Channel 12 are compared with previous values for Channel 12 in order to calculate a delta. It is the same for Channel 13:

Calculated Deltas

Server	Channel	Msgs Sent	Msgs Rcvd	...
Server1	12	77	249	...
Server1	13	11	34	...

To accomplish this on the client side requires that one must keep a “cache” of prior data as obtained in the previous time interval. Additionally, each row of data must be cached by one or more “index” columns that uniquely identify the source. In this case, it is a combination of the Server and Channel columns that identifies the rows to be compared.

This is not overly difficult, as there are simple algorithms for constructing a “key” from the index columns and storing a data row in a hashtable using that key. For each new row of data, old data are extracted using the key, a delta is computed, and the new data stored in place of the old.

However, it does put a burden on the developer of the monitoring client to maintain a cache and to calculate deltas against the incoming streams of data. Had the delta values been computed and exposed within the MBean, no such effort would be required.

In practice, it gets worse. Often, the number of sources does not remain constant. Channels may be created dynamically and the ID of each new channel is continually changing, starting at 1 and incrementing each time. Old channels may be closed and their IDs never reused.

In this situation, the cache that we maintain in order to perform the delta calculation keeps growing:

A Delta Cache That Keeps Growing

Server	Channel	Msgs Sent	Msgs Rcvd	...
Server1	12	12936	10566	...
Server1	13	9467	8789	...
Server1
Server1	3999	12341	7785	...

Here, channels 12 and 13 may be long gone, but the previous values stored in the cache are still there. To avoid a serious memory issue, the client code must provide a mechanism by which the cache can be cleared of items that are no longer needed. To do this requires that an event be generated, indicating that a channel has been closed.

Of course, we may need a “channel closed” event for other reasons, but to require it just so we can safely calculate deltas seems excessively burdensome.

Because this problem is seen so often, advanced visualization products usually provide built-in caching capability and transformations that can be used to perform the required work. However, when trying to apply a low-level charting package to the problem, one quickly finds that the problem is much bigger than originally thought.

The obvious recommendation here is to move the delta calculation back to the MBean code. In most cases, a “count” metric is going to be plotted, so provide the delta calculation up-front. It is usually much easier to do this at the source. There is no need for indexing of the data as the source contains both the old and new data. There is no need to worry about the coming and going of the objects (e.g. channels) as the delta processing is self-contained within the object.

This is not to minimize the problems on the data-collection side either. There are difficulties that arise when moving the job back to the server. Specifically, the problem now becomes how to manage the information so that it can be accessed by multiple clients without having to keep track of which client knows what.

The most common way to deal with this is to provide the delta in terms of a “rate” rather than an absolute delta. This way the time interval for the calculation can be maintained independent of the client. The client can be supplied both the

rate information and the original total information. In many cases, it is the rate that users want to see anyway, as the collection interval can vary and is really irrelevant.

RECOMMENDATION: Perform delta calculations and/or rate computations on the server side rather than on the client. It is much more efficient and easier on the client-side developers.

D. Give me a break... part 1

Another problem often seen in data collection systems is the inefficient encoding of information that is passed to the client for presentation. A good example of this is the “Statistics” column in our sample MBean:

String Encoded Data

Server	Chan	...	Statistics	...
Server1	12	...	Proc=124ms,wait=34ms	...
Server1	13	...	Proc=1196ms,wait=98ms	...

Often this is seen when the default view of the data is a simple HTML page. Having the metrics already available in a string form makes it easy to display in an HTML table without having to do numeric formatting.

Sometimes, developers are inclined to use XML since it is portable, easily parsed in client systems, and it is not binary:

```
<?xml version="1.0"?>
<Statistics>
  <ProcessTime units="ms">124</ProcessTime>
  <WaitTime units="ms">34</WaitTime>
</Statistics>
```

One argument for using either form is that the metrics encoded in the string are to be accessed as a unit, so it is clear that they represent a metric at the same instant in time.

However, most monitoring and visualization systems are fully capable of dealing with structured numeric data. JMX provides a simple OpenMBean data type called CompositeData. Several related measurements can easily be placed into a Composite structure using JMX API calls. The data are transferred via JMX in an efficient binary form.

With XML or string encoding, the server must pack the data into a String and the client has to unpack it at the other end, knowing the schema for the data. Using CompositeData, a client has information immediately available about the semantics of the data. With Strings, it has none.

RECOMMENDATION: Use structured Composite and Tabular data types wherever possible and avoid string encoding that results in extra network overhead and more development work on the client side.

E. Give me a break... part 2

A monitoring application is often required to “slice and dice” data it has collected, typically in the form of charts and graphs showing data “grouped by” various dimensions.

In our sample monitoring application the data that are associated with two sources: the Channel, and the Server on which that Channel is implemented. The same Channel may exist on different Servers.

A user may want to see the total messages sent/received across all Channels on a server (aggregation), or see how the traffic for all channels is distributed across different Servers (breakdown). Since the Server, Channel, and message counts are all available in the same MBean, this is not a problem. A Group By operation can be performed on the data, and the appropriate chart used to present the results.

However, look at what happens when the user tries to correlate the total message counts and the number of connections on each channel, using our sample MBeans. In order to perform a Group By operation relating Connections and Total Messages, the data need to be in the same table. It is necessary to perform a join operation on the two independent tables before the data can be used for this purpose.

Upon review, it can be seen that there is a one-to-one correspondence between the rows of data in the first table, the Channel table, and the rows of the second table, the Network table. Because the data structures inside the application were maintained separately, the JMX MBean was designed to expose the data in two different tables. It might have been better to combine these two tables ahead of time in order to minimize work that has to be done on the client side.

Joined Channel Data

Server	Chan	Msgs Sent	Msgs Rcvd	...	Stats	Conns
Server1	12	12549	9613	
Server1	13	9456	8719	...		

It is not suggested here that you build analytics into your MBeans. Rather, collapse the data structures to their minimal form before exposing the data. One would not say that the answer to an algebra problem was $2x + 3x + 5 + 9$. Instead, it would be simplified by combining common terms and providing $5x + 14$ as the answer.

RECOMMENDATION: When there is a one-to-one relationship between available metrics, combine these into a single bean to avoid having to perform Joins and Combines on the client side.

F. Have some foresight...

In the sample ChannelInfo MBean, memory metrics are stored in a 32-bit integer, assuming the common Java limitation of 1 GB heap space.

The integer representation works fine as long as you are only looking at one bean at a time, as in the simple HTML interface to the beans. However, when aggregating this metric across multiple beans, we get the following:

Totals for Max Bytes

Server	Chan	...	Max Bytes	...
Server1	12	...	536,870,912	...
Server1	13	...	536,870,912	...
Server1	14	...	536,870,912	...
Server2	11	...	536,870,912	...
Server2	13	...	536,870,912	...
Server2	14	...	536,870,912	...
...
Total			3,221,225,472	

The total for all the channels is a number that is well over the largest positive integer – 2,147,483,648 – which can be represented in 32 bits. The integer worked fine for one channel, but not multiple channels. The data type should have been long from the beginning, anticipating aggregation.

RECOMMENDATION: Look ahead to the results of aggregations and use data types that are appropriate. Aggregate totals can quickly grow very large.

G. Tracking History ...

A common requirement for monitoring systems is to maintain a historical record of activity. One solution is for an MBean to write to a log file and record its own history. However, in real world systems, where data volumes may be huge, this is not practical. Storing the data into a relational database provides the user many more options for reporting than would simple log files.

Often overlooked is the need to include a timestamp with the data. This is easy to do of course, but suffers from a major problem. Activity on the network can delay acquisition of metrics data. If a timestamp is assigned at the time data are received in the client, any computation of rates or averages that makes use of the delayed timestamp may be inaccurate.

RECOMMENDATION: Include a timestamp with all data intended for archival or time-based analysis. It should be stored at the time of data acquisition and be in millisecond resolution in order to provide the most precise calculations.

H. To poll or not to poll...

Given the large amount of real time data that can be produced in a monitoring system, one question often comes up: can the notification capabilities of JMX be used to minimize network traffic and overall load on the system?

Typically, a monitoring system is initially developed in a polling mode. On a regular time interval, e.g. 10 seconds, a request is made to query various metrics and the data are transmitted back to a client system for analysis and display. This polled approach can be costly in terms of network bandwidth and processing overhead.

However, looking to notifications to solve this problem may be somewhat fruitless. Notifications do not provide a panacea for all that is wrong with a monitoring system. In fact, if used incorrectly, there could be even more overhead introduced.

The requirement to store historical data means that metrics must be obtained on a regular basis, whether or not anyone is looking at them. For example, the Total Message counts and Bytes Used metrics are not candidates for notifications. These must be polled in order to maintain a consistent history of the values. There is nothing to be gained by using notifications.

On the other hand, the Connections count may not be changing every 10 seconds. A connection could remain alive for hours or days. In this case, the use of notifications could reduce the network traffic by only sending data about the connection count when it changes.

A number of other issues should be considered when using notifications. These are often overlooked, yet require development support to properly use notifications:

1) The use of notifications must be combined with polling and/or a caching mechanism. A notification is not issued until data changes. When a display page is brought up, showing the current count of connections, it will initially be blank and will fill with data only as the Connections are added or removed.

The client application needs to populate a table or chart with the current set of values either by polling for that data on display activation, or by using a cache that maintains current table values independent of the active displays. This functionality requires that equivalent attributes be provided for any values obtained via notifications.

2) When used in conjunction with historical data obtained from an archival database, the use of notifications involves similar complexity. A trend chart when first brought up must be populated with data obtained from the archive. As notifications are received, those values must be appended to the chart. The archived data must be requested only once.

Building the mechanisms to support the merging of notified data with current or historical data can involve a fair amount of development effort, often minimized in the early discussions about building a monitoring client against newly minted JMX data.

RECOMMENDATION: Use notifications where data are not changing regularly and there may be some real reduction in overhead. Do not use notifications for everything just because they are available. Design MBeans to support the integration of notified data with current or historical data.

IV. SCALABILITY AND MAINTENANCE

There are many more pitfalls that may occur with monitoring using JMX. Most important of these are in scalability and maintainability. Some systems simply produce too many beans and this can result in terrible performance. In others, the complexity of the MBean names and key properties is overwhelming and can be a huge maintenance burden.

There needs to be the usual tradeoff made in balancing complexity against performance. Many systems return one row of data for each MBean, using the bean name to encode the source. This can get unwieldy if overused, and the number of beans can grow excessively. An alternative might be to design beans that return multiple rows of data, e.g., instead of an MBean for each channel, provide an MBean for the server and return a table containing a row for each channel.

Other issues come into play when one looks at common usage patterns of the MBeans in a system being monitored. Are the data polled randomly or is there a predictable sequence that can be used to pre-fetch data and have it available on the next cycle?

For maintainability, it is important that data formats for the beans remain stable and are not changed without some sort of upward compatibility plan and/or deprecation plan.

For the most part, a combination of best practice knowledge and good common sense can help produce a quality monitoring system that performs well and has all required functionality.

REFERENCES

- [1] SUN – Java Management Extensions (JMX) - Best Practices, 2007
- [2] Justin Murray, HP - Design Patterns for JMX and Application Manageability, October 2004.
- [3] BEA – WebLogic Server – Developing Custom Management Utilities with Version 9.2, JMX, February 22, 2007.
- [4] SUN - J2EE Management Specification JSR-77, Java Community Process.
- [5] Peltz, Kumar - Apply JMX Best Practices, Java Pro, December 2004.