



P i r e a n

Measuring the availability and performance of business processes

By Mike Cartwright, Pirean Technology Solutions.
Version 1.2



Table of Contents

1	An introduction to availability and performance	4
	Business service management	4
	Application level management	4
2	Monitoring business processes	5
2.1	The users perspective	5
2.2	Simulating a user experience	6
2.3	Key benefits	6
3	Tools for monitoring performance and availability	7
3.1	Application Response Measurement (ARM)	7
3.1.1	Embedding ARM	7
3.1.2	ARM Simulation	7
3.2	IBM Rational Robot	8
3.3	IBM Tivoli Manager for Transaction Performance (ITMTP)	8
4	Case Study – Manufacturing Company	9
4.1	The requirement	9
4.2	The solution	9
4.3	The real-time interface	9
5	Conclusion	10
5.1	About Pirean	10



“The rising tide of international standards and legislation focuses expectations ever higher, forcing business to improve internalised standards and deliver service assurance. In this environment your organisation cannot allow technology to dictate your service path or obstruct progress. If it does, then it isn’t working and neither is your business“

An introduction to Utility IT, Pirean Limited.



1 An introduction to availability and performance

In today's on demand world your IT infrastructure plays an ever increasing role in the support of business processes and delivery of products and services.

Given the increasing complexity of technology infrastructures and the demand for new systems, services and access points (sharing data across your business, with your customers and partners), it's important to consider how you deliver a stable service.

Part of Pirean's Utility IT model, our advanced infrastructures portfolio facilitates the proactive management of your IT systems, not just at the server or network level, but from an end-user perspective through a real-time monitoring model.

Business service management

For many businesses today the question isn't whether your servers are available, but whether your key business processes are online and meeting demand.

When you consider the complex mixture of machines, applications, network traffic and user-interactions which make up a business service you have to ask yourself, "How do you know that the applications supporting your critical business processes are available?"

In many organisations the answer is simple: if all of your servers are up and your applications are running; if there are no known errors and the network has not reported any faults – then it's a good day.

But what you are *really* saying is that you know that all monitored components are available; and with every assumption you make, you distance yourself one step further from monitoring what's really happening in your business.

This approach to availability and performance management doesn't ensure your services are meeting your target SLAs or qualify the experience of the end user or (more importantly) your customers. What's really happening is that you're *assuming* that if all of the components that are important to your business are monitored then the system *should* be working.

Application level management

Applications today have more dependencies across more systems, often distributed over multiple operating systems and spread over an ever wider geographical area than traditional client/server implementations. Functions are partitioned over networks, solutions exploit many different technologies - complex availability and load balancing solutions make it difficult to assess the true business impact of a failure.

The solution here is to look to alternative monitoring solutions; in many cases you can provide a 'real' monitoring solution through simulation. By simulating a user running key transactions and monitoring the performance of each process or transaction, you can constantly **prove** the availability of your IT solutions that support your business processes.



2 Monitoring business processes

In order to keep people focused on business, your technology infrastructure needs not only to work, but also clear the way to realising your company's true potential.

From planning the most efficient use of a factory to providing customer details on demand, companies rely on technology to facilitate their business processes. It is this recognition that leads us to monitor our infrastructure. Monitoring aimed at giving us the confidence that the systems critical to our business are available.

Traditional Monitoring focuses on individual components, ensuring that servers are available and applications are running. But this approach does not naturally accommodate complex relationships between components, or indeed an understanding of which components need to be available and performing in order for the end-user to work successfully and without interruption.

It is always possible that there is a component in the end-to-end solution that is not being monitored. As the granularity of monitoring solutions increase the volume of events to be filtered and understood increases – making it proportionately more difficult to monitor a real service over individual components.

It's also important to note that even if every component is functional, a system which is under-performing is of little more use to the end-user than a system that is not working at all.

2.1 The users perspective

Consider a user who works in a call centre, talking to hundreds of potential customers each day. This user accesses an application that provides details on your companies' products.

Behind the scenes there are many components; servers, networks, backend systems and a database which are all supporting the end-user application.

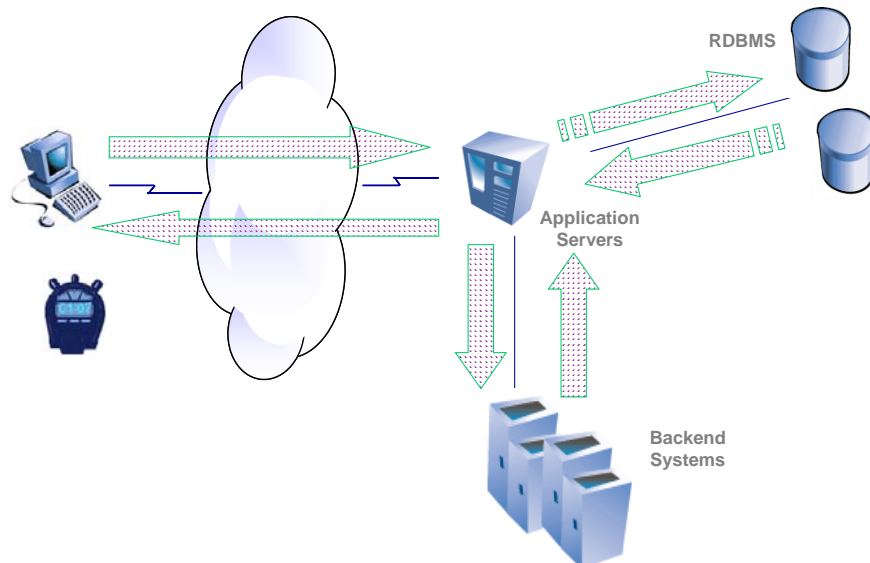


Figure 1 – Example of systems supporting a user's application



As in any such instance, the user does not, and should not; have to understand the technology behind the scenes. They just need the applications they use in their day-to-day role to be available and responsive.

2.2 Simulating a user experience

If a workstation in your call centre was simulating a user by constantly running the transactions that form business processes, you would be constantly **proving** the availability of the IT solutions that support the processes.

With this approach the monitoring model becomes very simple - as you no longer need to think about individual components but the end-to-end solution.

In this example, a simulated user in the call centre could:

- Connect to the application
- Log on
- Perform transactions
- Confirm that the performance is within acceptable service levels

To provide a complete solution, it is still recommended that traditional monitoring is implemented to act as a pro-active tool on each unique component in order to identify potential issues. However we now have the added reassurance proving that the business process is working from the end-to-end perspective.

2.3 Key benefits

End to End availability and performance monitoring provides three advantages over traditional monitoring:

- Availability
 - You can prove the availability of your IT solutions that support your business
- Service Level Monitoring
 - The response times prove that the solution is working and performance is within agreed service levels
- Capacity Planning
 - You can see the impact of the load on the systems at different times of the day
 - Performance data can be analysed to detect degradation of performance before the user is impacted
 - Data can be extrapolated to plan for systems growth



3 Tools for monitoring performance and availability

There are several tools available for monitoring performance and availability, most utilize the ARM standard.

3.1 Application Response Measurement (ARM)

Developed by The Open Group, The Application Response Measurement (ARM) standard describes a common method for integrating enterprise applications as manageable entities. Through the provision of an application programming interface (API) the monitoring of availability and performance data is possible from within and across diverse applications and systems.

In simple terms, an application starts an ARM timer prior to performing a transaction, and stops the timer just after the transaction completes. Essentially starting and stopping a stopwatch. An agent on the workstation running the application measures and monitors the transactions, recording and making the data available to management tools.

ARM calls can be nested to handle more complex processes, for example:

Complete Transaction - Start timer

Perform actions

Transaction A – Start timer

Perform actions

Transaction A – Stop timer

Perform actions

Transaction B – Start timer

Perform actions

Transaction C – Start timer

Perform actions

Transaction C – Stop timer

Transaction B – Stop timer

Perform actions

Complete Transaction - Stop timer

3.1.1 Embedding ARM

If the application to be measured is one that is developed in house, the ARM calls can be embedded into the application.

This has a number of benefits; firstly there is no need for a dedicated workstation to simulate transactions, real users are performing real transactions every day. This information can be collected and analyzed without any impact on the day-to-day operations of your business. Secondly, adding the ARM calls to custom code allows a very detailed level of data to be collected. Within the application you know exactly what is happening, so you can measure time for a page to be displayed, time for data to be received from database, time taken for application backend to process data and the complete end to end transaction time.

3.1.2 ARM Simulation

If the applications being used are not developed in house, or the work to revisit the code and add ARM calls is deemed too great, then an alternate option is user simulation.

In this instance, a dedicated workstation would be placed in the location from which you wish to measure Availability and Performance.



The robot is configured to perform one or more transactions, captured in a script generated using a tool such as IBM Rational Robot.

During the playback, the robot will make arm calls before and after each part of the transactions being measured according to the design of the script. Within the script you can see each transaction as a series of key presses and mouse clicks, ARM timers can be added at any point.

There could be many identical Robots placed throughout your organization to measure the same transaction from different locations.

3.2 IBM Rational Robot

There are a number of tools available to enable the recording and playback of transaction.

The IBM Rational Robot product is essentially a GUI-based record and playback tool. The Robot software provides the functionality required to automate transactions using any Windows-based application. It assists in the creation of a script by allowing a user to record a transaction as they perform it. This script can include window verification points and other features to ensure the activity was successful.

The resulting script is Visual Basic making it easy to customise and extend.

An unattended playback can be performed on any workstation with the Robot Software installed.

3.3 IBM Tivoli Manager for Transaction Performance (ITMTP)

The ITMTP product supports any application that has embedded ARM calls, and controls the collection of the data generated. The data is returned to a central database for analysis and reporting.

ITMTP fully integrates with the suite of Tivoli Enterprise monitoring solutions, meaning that monitors can raise events and where configured tickets based on thresholds or when the application becomes unavailable.

In the scenario where many robots are placed around a company's environment, Tivoli can assist in the correlation and intelligent interpretation of events.

For example:

Site A is reporting that the Application is unavailable.

Site B is running transactions successfully

Site C is running transactions successfully

A Network Device on Site A is reporting an error

Intelligent correlation allows us to determine that the application unavailable error from site A can be ignored as the application is available from the other sites. The event can be correlated with the Network Device error, which is the cause of the problem.



4 Case Study – Manufacturing Company

4.1 The requirement

A large manufacturing company implemented an Availability and Performance solution in 2004, aimed at measuring the response time and availability of its SAP solution, worldwide.

There was a perception from users that new SAP systems and the relocation of data centers had resulted in performance issues. In truth there were many factors associated with the new solution, new interfaces, new procedures, new processes each of which had an impact.

Without performance data there was no way of measuring the response times. Confidence was decreasing in the new solution, without any evidence to show that there was any real issue.

The company recognized a need to measure availability and performance from the users location, and to publish the information.

4.2 The solution

The IBM Rational Robot tool was chosen to simulate user activity on each of the different SAP servers (R3 Server, Activity Planner, Enterprise Buyer Professional, Customer Relationship Manager, Business Information Warehouse and Workplace)

Robots were placed on sites throughout the world, from England to Sydney. Every 10 minutes, each workstation logs onto each SAP server, recording the availability and response time for the chosen transaction.

IBM Tivoli Monitoring for Transaction Performance was chosen to collect and correlate the performance data. The tool was the natural choice, with an enterprise wide Tivoli implementation.

4.3 The real-time interface

The IBM tools allowed the data to be collected and stored in a central database. This would have provided a perfect repository for historical reporting, but the customer wanted to publish near real-time data on the intranet. By publishing the statistics users and managers could see that the systems were performing.

A web based application was developed to graph and display the data. Based around the simple concept of traffic lights, the first page shows a map of the world and a coloured light for each region. Clicking on a geographical region, for example Europe, gives a view of all markets being monitored in Europe. Clicking on a market will show each of the sites monitored within the market. Clicking on a site displays near real-time performance graphs for each of the systems being monitored from the chosen site.

The data collected can be presented in many different ways. The performance graphs for the SAP systems from each site are published on the same page, allowing a direct comparison of performance by site.

The solution has proved very successful for the customer, and is being extended to cover other applications in their environment. It is a key tool in maintaining service levels and has proven itself as an accurate and pro-active monitoring tool.

The solution is being extended to account for scheduled outages and maintenance windows and now forms a critical part of service level reporting.



5 Conclusion

Through end to End availability and performance monitoring:

- You can prove that a business critical application is available
- You can publish near real-time response times that show Service Levels are being achieved
- You can see the impact of system load at different times of the day
- You can analyse Performance data to detect any degradation before the user is impacted, providing pro-active monitoring of your business processes.

5.1 About Pirean

Focused on adaptive enterprise management, Pirean's solutions portfolio provides a best-practice foundation of enterprise availability, provisioning, orchestration and business service management.

Built on IBM Tivoli and Microsoft software platforms, our solutions align you infrastructure management programme to your business, encompassing business service and system level visualisation.

As part of the Utility IT model Pirean's systems management portfolio delivers a complete, end to end solution – from service realisation, monitoring and configuration management to security profiling, compliance and identity management.

Bringing together Systems, Service and Security management facilitates advanced automation and relationship management, providing the framework you need for an on-demand infrastructure.

For more information on Pirean's Utility IT solutions visit www.pirean.com



P i r e a n
www.pirean.com