End-to-End Monitoring for the TIBCO-powered Enterprise

Becoming More Proactive through Centralized Monitoring and Alerting

Introduction

Integration middleware from platforms such as TIBCO, Apache Kafka, and Solace represent unique monitoring challenges for technology support teams. Since these teams have specific domain knowledge and specialization, monitoring and alerting tend to be siloed and challenging to share with Management and across workgroups. Gaining visibility and context across these middleware-specific silos and workgroups can be achieved with SL's RTView® Enterprise Edition which enables improved communication and collaboration between these groups in a scalable and efficient fashion.

For most organizations, the greatest return on investment in monitoring comes from preventing outages and degradations in the first place. Organizing and presenting information to users in the context of the applications is the key to improving communication between key workgroups and becoming more proactive. It also increases alignment with the business. There are plenty of tools available to troubleshoot problems once they occur, but these can come at a cost in productivity and downtime. Especially since time wasted understanding the dependencies and re-engineering the app means significantly increased MTTR (Mean Time to Resolve an issue).

This document describes the four key attributes that RTView Enterprise brings to provide a global perspective to monitoring in your organization. Each is presented in the context of a common IT problem, how to best solve it, and the mechanism by which RTView® Enterprise addresses that problem.

Organizations typically deploy end-to-end monitoring for two primary and quantifiable objectives:

- Reduce application and business service outages
- Reduce the time it takes to restore an application into production after an outage or degradation occurs (MTTR)

Several impediments are often encountered when implementing commonly-used middleware monitoring solutions. We will explore each problem area in more detail.

1. Monitoring silos limit ability to share information
2. Inability to correlate performance across complex, distributed applications
3. Alert overloads result in obscured visibility
4. Difficulty Understanding How the Application, as a whole, is Functioning

PROBLEM 1

Monitoring Silos Limit Ability to Share Information

In large organizations, distinct workgroups are responsible for different parts of the technology stack. Infrastructure, middleware, or application groups collect volumes of monitoring data about important subsystems. But this information is typically available only within the group and is not visible to others in the organization. These groups are often referred to as “silos” because of their isolation from one another. Since business-critical applications and services are built on top of multiple technologies, the lack of a common monitoring and alerting framework means technology support teams only have part of the picture and can become stuck in a reactive mode.

Solution

Consolidate Monitoring Data Across Technology Silos with a Unified Monitoring Architecture
RTView Enterprise provides an architectural backbone and service model to support a consolidated view across the multiple distributed technologies supported by these independent workgroups, thus addressing the silo problem. A standard, unified and extensible console provides access to important data collected by other monitoring tools, including TIBCO RTView stand-alone monitoring products.

Centralized service alerting provides early warning and enables support teams to receive alert notifications, identify, and address the cause of the problem before the service or application is affected.

Benefit

Better Visibility Across Support Teams
Each group retains control of its data yet can expose the data to a wider audience, improving communication between teams. The result is proactive avoidance of problems, rather than the typical reactive approach to troubleshooting that occurs when one group is not aware of issues occurring in another part of the stack.

Example

Monitoring and diagnosing an application powered by TIBCO BusinessWorks, Oracle databases and VMware VMs may employ the management suites from each environment (TIBCO Hawk, Oracle Enterprise Manager, VMware vCenter). With each of these suites focused more on managing than monitoring their components, development and support teams struggle to get the visibility they need for their critical applications.

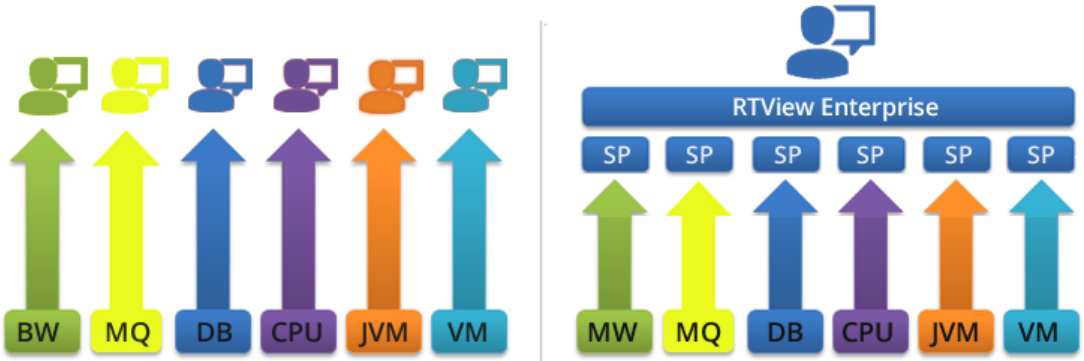

PROBLEM 2

Inability to Correlate Performance Across Complex, Distributed Applications

Modern complex applications and services are often implemented across multiple load-balanced middleware components. Support teams must access data from several servers, one at a time, to determine aggregate throughput of the system. Because data from different technologies are collected in different locations, it is also necessary to access disparate monitoring tools to investigate a problem. This makes it challenging to be proactive in preventing them and time consuming to troubleshoot them.

Solution

Aggregate and Correlate Key Metrics Across Heterogeneous Components

RTView Enterprise provides the framework to correlate Key Metrics across all components on which applications depend, regardless of their location. It can aggregate and total metrics across load-balanced servers and provide end-to-end visualization of complete transaction paths across technologies. This makes it possible to understand causal relationships between different technology components and applies to both current and historical data.

Benefit

Early Warning, Rapid Recovery

With service-based notifications and all data readily available in a single console, support teams do not need to waste time or be trained to search for important metrics in different tools. This is especially important during stressful outage events. The Key Metrics capability makes it possible to visualize the overall health state of all components affecting an application well before alerts occur – providing early warning that there may be degradation in the system. Users can avoid problems altogether rather than troubleshoot them later.

Example

Running multiple TIBCO BW instances in a VMware private cloud environment. A rogue process from an unrelated application on the same physical host begins to consume system resources and brings the server to its knees. Without the ability to correlate host, VM and TIBCO performance and health state metrics, three separate support teams using three different monitoring consoles may significantly delay problem resolution.

A history heatmap enables correlation of key metrics between the different technologies in a service

PROBLEM 3

Alert Overloads Result in Obscured Visibility

Support teams often struggle with alert fatigue and alert overload with too many alerts cascading in from too many tools. The problem can be so pervasive that for most organizations, one or more alerting systems are ignored. And this problem is exacerbated during outages when rapid resolution is most critical.

Solution

Service-centric Alerting

Service-specific alerting is an efficient way to channel alerts to support teams. By providing filters for one or more:

- specific services
- specific technologies such as TIBCO BusinessWorks
- environments such as Production
- areas such as a Data Center
- Groups
- Owner

. . . support teams only receive the alerts they care about. A flood of alerts rapidly becomes more manageable.

Additionally, different services can be ranked with varying levels of business criticality to ensure the most critical services are prioritized above less important ones.

Benefit

Operations and Support teams can maintain focus on the services that matter most without the distraction and resulting alert fatigue that comes from unfiltered alerting.

Example

A support engineer receives a service alert notification on his Smartphone for a critical business service and is immediately able to drill down to see an alert list for the complete service. He can then drilldown to the alerting component, understand if the alert is a trend or a spike, and take remedial action.

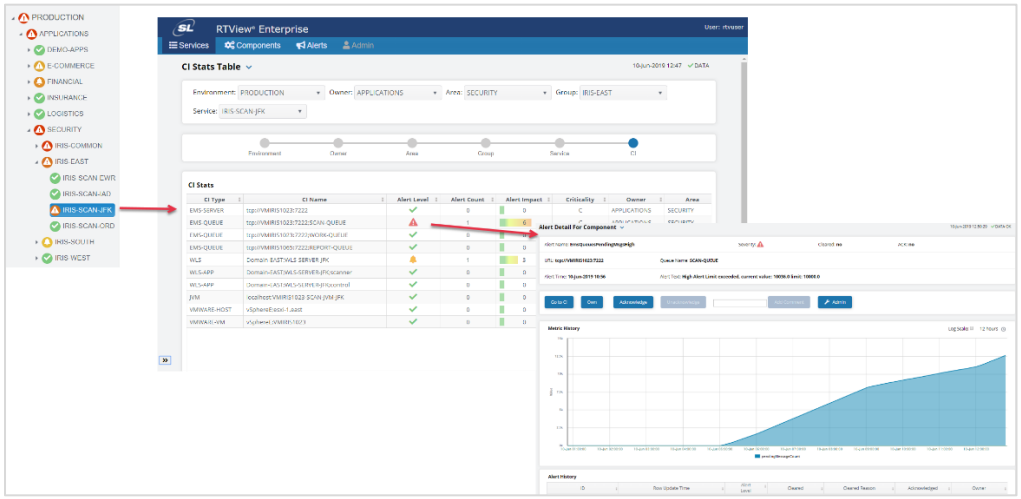

Service alerts enable support teams to understand the business impact of a component issue on the entire service and drill down to alert detail and history

PROBLEM 4

Difficulty Understanding How the Application, as a Whole, is Functioning

Many popular monitoring tools, including Application Performance Management (APM) tools, do not provide visibility into the source of middleware issues and the impact those issues have on critical business services. These tools are good at addressing certain types of problems, especially for identifying and addressing code level issues and for use in non-production environments. However, in production, operations teams need to answer higher-level questions about application health and performance which can be difficult to do with most monitoring tools.

As an example, APM tools provide an external look at the JVM and not deep visibility into the TIBCO middleware running inside the JVM. Only looking at the JVM externally can be a significant limitation because most TIBCO BusinessWorks problems occur at the activity or process level and not at the JVM level. Moreover, an APM tool can tell you where a transaction failed but not why since they have no transactional visibility into the entire TIBCO domain.

Solution

Application-centric middleware monitoring is designed from the ground up to provide visibility into the middleware which is often opaque to other types of monitoring approaches. RTView Enterprise pulls monitoring data from TIBCO monitoring APIs such as the EMS Admin, JMX, and TIBCO Hawk (without deploying new agents and the associated management complexity and maintenance costs).

Benefit

Support teams often find that application-centric tools for monitoring their services and supporting technologies in production can be more cost-effective, agentless, and lower overhead monitoring approach compared with other tools.

Example

In an extensive TIBCO application environment where there are hundreds or thousands of instances of TIBCO BW or EMS, operations teams can waste a lot of time determining which component supports which application. Therefore, understanding business impact, prioritizing support efforts, and diagnosing how a particular instance is affecting application performance or end-user experience is easily accomplished.

A heatmap shows the real-time health of different components supporting services across the entire estate

SL is a San Francisco Bay Area-based company that provides End-to-End Application Monitoring and Middleware Monitoring for Global 1000 and mid-market companies that depend on custom, high-performance applications. Over the past 25 years SL's exclusive focus on real-time monitoring and visualization, commitment to customer success, and partner-centric culture are why thousands of industry leaders have chosen to work with SL to support their most critical applications and businesses. SL's RTView® product lines address a broad spectrum of enterprise visibility challenges spanning application and service availability, performance monitoring, component-level infrastructure monitoring and custom monitoring.

For more information on SL, please call +1 415-927-8400 or visit www.sl.com.

SL Corporation | 240 Tamal Vista Blvd. | Corte Madera, CA 94925 | USA
+1 415.927.8400 phone | +1 415.927.8401 fax | www.sl.com