

Combining Technology and Business Centric Monitoring to Minimize Customer Incidents.

Vantify Experience Centre Whitepaper

Business-centric monitoring based on real-time analysis of ‘Business Events’ complements technology-centric monitoring to provide a complete and real-time view of business and service performance. The use of Complex Event Processing (CEP) technology as the ‘engine’ for business-centric monitoring further enriches insight into service behaviour enabling earlier incident detection, better customer experience and improved technology Return on Investment. Incident detection rates can be improved by up to 80% leading to substantially reduced operational costs. This Whitepaper contrasts each approach, arguing that an effective service or activity monitoring solution must combine elements of each to maximise business performance, operational efficiency and customer satisfaction.

Technology and Business Events

This whitepaper considers two complementary approaches to monitoring the performance and behaviour of business activities and service delivery for IT intensive enterprises, namely **Technology-centric** and **Business-centric**¹. The defining characteristic of each approach is the nature of the events upon which monitoring, analysis and alerting is based. Business-centric solutions, such as WestGlobal’s Vantify Experience Centre, are driven by events with real business significance, hereafter termed **‘Business Events’**. Solutions of this nature are typically referred to under the banner of ‘Business Activity Monitoring (BAM)’. Technology-centric solutions are based on monitoring lower-level technical events for domains or individual elements of infrastructure (e.g. networks and servers); this paper refers to these as **‘Technology Events’**. Examples of each event type are provided in Tables 1 and 2 below.

Business Events
Target of 650 Postpaid customers provisioned since 8am has been met.
25% of pending number porting requests have exceeded the SLA target of 2 hours.
User request to top-up / recharge prepaid account received at 8:31:03 (hh:mm:ss).
SMS average delivery time increased by 200% for messages to Operator X in France.

¹ For a more granular or detailed breakdown of monitoring approaches please refer to our Whitepaper ‘Total Insight Customer Experience Management [WG/WP/02-2008-01]’

Broadband activations for last 24 hours are 1. 27% below forecast 2. 31% above same period last week
Zero activity for content downloads from Content Partner 'Acme Ltd.' in last 60 min.

Table 1 - Business Events

Technology Events
CPU utilization exceeds 85% on Server42.
37200 HTTP requests to www.westglobal.com
Disk failure in redundant array DA07.
Server room humidity alarm.
WAP Portal response time has breached 3 seconds.
Garbage collection process consuming 95% of CPU.

Table 2 - Technology Events

There can be overlap between Business and Technology events where a single event can be interpreted as either. For example, a WEB service allowing a user to check their prepaid balance may be interpreted as a Business Event but can equally be considered as a technology event, specifically an HTTP protocol request response to a WEB application server on port 80. As a general rule of thumb, however, business events have significantly higher semantic content and complexity compared to technology events. Typically a Business Event exhibits one or more of the following characteristics:-

1. It is associated with a recognizable business activity.
2. It has rich semantic content relating to a transaction or a customer.
3. It is technology agnostic.
4. It is complex and is often derived from other events.

As an example, consider the common service offered by mobile operators to send a notification SMS when a caller leaves a voicemail. We could define a single Business Event related to the successful operation of this business activity. It is complex as it based on aggregation of multiple underlying events (detection of the new voicemail, detection of the associated SMS request, potentially even detection of the time for the user to retrieve the voicemail). It is clearly associated with a business activity rather than behaviour of an underlying technology component. Finally it has high semantic content as it contains data related to a customer and the multiple elements of a transaction. This data may be extremely rich relating to not only success or failure of the transaction, but also containing timing information, user identifiers and other contextual parameters.

Sources of Incidents

A Gartner study² has shown that approximately 20% of IT-related service incidents are due to hardware or infrastructure problems - see Figure 1 below. As hardware reliability continues to improve and reducing cost makes redundant or replicated design increasingly economic, it can be expected that this percentage will reduce even further in the future.

It follows, therefore, that a technology-centric monitoring solution will have a maximum incident detection efficiency of around 20%. The remaining 80% of incidents can be largely attributed to software bugs, interworking, configuration and maintenance issues. These incidents frequently lead not to immediate failure but to a degradation in performance or increase in errors that are nevertheless severely service impacting and much harder to diagnose compared with those related to infrastructure failure. In fact, these incidents can only be reliably detected and diagnosed through an approach that monitors complete service behaviour, measures performance across all relevant business metrics and is able to detect any degradation that will impact on customer experience. The class of solutions providing this capability are generally referred to as **Business Activity Monitoring (BAM)**.

An important conclusion is that a complete and effective enterprise monitoring solution will combine elements of technology and business -centric approaches to ensure the maximum possible incident detection efficiency. The next section of this paper considers the characteristics of each approach in more detail.

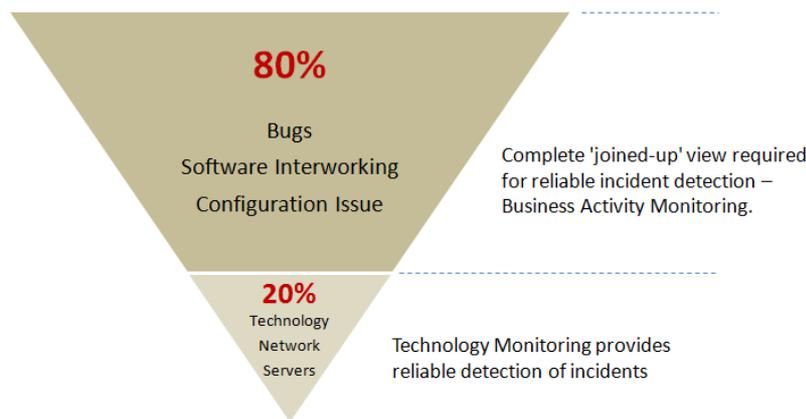


Figure 1 - Approximate Distribution of Incidents by Cause

² Gartner Group Research Note.

Contrasting Technology and Business -centric Monitoring

This section considers the characteristics of technology and business-centric monitoring in more detail. The concept of a **Service Model** is also introduced - this provides a framework against which business and technology events may be mapped onto customer facing business activities. The mechanism by which the service model is configured and maintained is a key driver of the value and insight provided by a monitoring solution. As shown in Figure 2 below, both technology and business-centric approaches can support a service model but, as this paper will show, there are important limitations associated with the service model for technology-centric solutions.

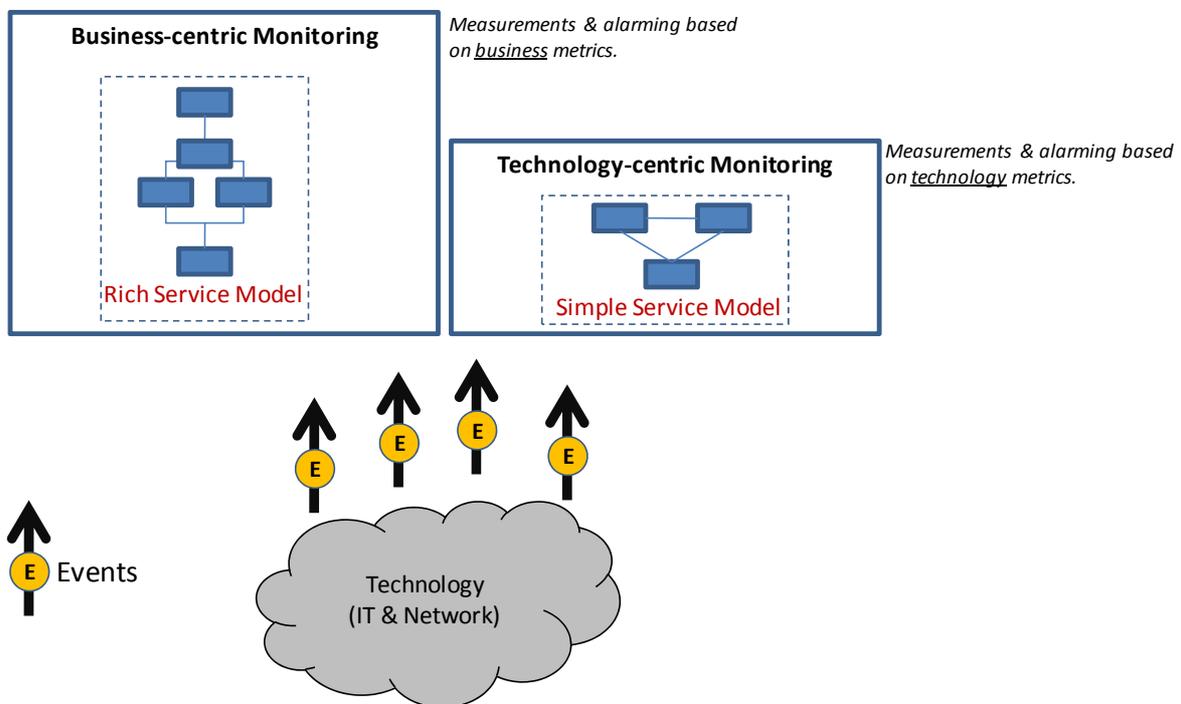


Figure 2 - Business & Technology Centric Monitoring

Technology-centric Monitoring

Technology-centric monitoring focuses on the measurement and alerting of indicators associated with the operation of a technical node (e.g. server) or domain (e.g. network). Events detected by such solutions are typically indications of failure or simple thresholds, for example an application server CPU utilization exceeding a configurable value.

Technology monitoring often generates huge volumes of data or events, much of which has limited relevance in alerting actual or potential incidents. It can be a non-trivial and very resource intensive task for operations staff to analyze this data to identify abnormalities in performance, particularly if this needs to be done in real time. For example, setting an alarm on CPU load exceeding 90% may be

inappropriate while an intense calculation is being performed, but valid under lower load. Without the context of the business activity, many technology events lose their significance.

Most importantly, technology monitoring often indicates that all network and technology components are operating correctly but performance problems are nevertheless being experienced with the business activity or service. This aspect is perhaps the single most compelling argument for complementing technology centric monitoring with the business centric approaches described later in this paper.

In terms of defining a Service Model, technology centric monitoring broadly offers two approaches, namely **Topology Discovery** and **Static Mapping**.

- **Topology Discovery** uses strategically located sensors to report back protocol activity to a central application server. By monitoring standardized fields in protocol headers, for example Port ID and Protocol Identifiers in the IP protocol, it is possible to reliably detect a range of services operating between different nodes. For example, a server receiving TCP traffic (IP protocol ID = 6) to Port 80 is almost certainly acting as a WEB server. Similar approaches can identify flows related to Oracle, SAP and a range of other services and products. This information can be used to dynamically build and maintain a topology map or service catalogue as part of a Configuration Management System. By monitoring and recording a history of the typical traffic profiles for each of these services, it is also possible to detect and alarm against unusual behaviour.
- **Static Mapping** is based on meta-data that defines dependencies between underlying technology infrastructure and the business activities dependent on that technology. This enables a technology-centric monitoring solution to indicate a list of potentially affected services when problems are detected with the monitored technology. This approach is also referred to as **Business Service Management (BSM)**.

Technology-centric monitoring is an important element of an operations department's armoury to provide assurance around the health of underlying infrastructure but, as indicated in Figure 1, it is important to recognize these solutions will typically detect a maximum of approximately 20% of incidents. The use of techniques such as Topology Discovery and Static Mapping can provide some improvement but because these approaches continue to rely solely on technology events, they are unable to detect issues unrelated to infrastructure. In order to provide a substantial improvement and address the 80% part of Figure 1, it is necessary to deploy business-centric approaches as described in the next section.

Business-centric Monitoring

The defining characteristic of a Business-centric monitoring solution is the ability to monitor and alert in real-time against metrics that are **relevant from business and customer experience viewpoints**. For example, it should be possible to detect and alert on conditions such as:-

- It takes longer than 3 minutes to activate a new postpaid user from a Point of Sale (shop).
- The number of UMTS Call drops has increased by 20% over 1 hour.
- Rate of connection failures for content partners connecting to an SMSC has exceeded 5% growth within the past 5 minutes.
- The response time for customers to retrieve invoices via the Customer Portal has increased by 50% over the last month.

This level of insight is valuable for two reasons in particular. Firstly, it is fully aligned with business objectives and metrics in terms of actual Customer Experience. Secondly, degradation in service performance will be measured directly and immediately - it does not need to be 'inferred' based on the performance of underlying technology. This is important given Figure 1 that suggests the majority of performance problems or incidents are not related to simple technology failure.

It is clear that an important feature of business-centric monitoring solutions is that they are driven by 'business events' as defined earlier in this paper. To illustrate, consider a prepaid recharge operation. Part of this business activity is the monetary adjustment to the user's prepaid account. A technology monitoring solution would measure this step of the activity typically in terms of metrics such as volume of requests, response time, CPU load on the prepaid server. In contrast, a business-centric solution would see this event as a 'Top-up request of amount x Euros for user MSISDN'. **The ability to extract this level of semantic information from underlying events is a critical feature of business-centric monitoring solutions.**

It was mentioned earlier that this class of solutions are generally described under the banner of Business Activity Monitoring (BAM), a term originated by Roy Schulte of the Gartner Group. Additional features that significantly enhance the value of BAM solutions are **Behaviour Analytics** and **Drill-down Diagnosis**, each of which are considered below.

Behaviour Analytics allows the detection of patterns of business activity or service performance that can provide more effective detection of incidents than simple threshold monitoring. Using the example of a prepaid recharge, a simple threshold alarm could be configured to detect any failure related to a user recharge attempt. With a pattern based approach, it is possible to refine this alarm to trigger only if a user has, for example, 3 failures in a certain time period. This would eliminate the occasional failures that may be due to user errors while providing a more reliable indication of real operational issues with the recharge business activity. Another interesting example is to detect strange usage patterns (very low usage, many short calls/sessions) following a short period after account activation - this could point to problems with the activation of the customer's account or their ability to configure or use a service, either of which is a risk of 'fast churn' and loss of the customer. The decision to adopt Complex Event Processing technology as the underlying engine for

WestGlobal’s Vantify Experience centre allows the real-time detection of such patterns.

Drill-down Diagnosis provides the ability to investigate and diagnose the root cause for performance issues by viewing the state of the IT services and even individual technical elements underlying a business activity or service. It is important to note that this is fundamentally different from technology-centric monitoring as the process has been triggered by a degradation in a business relevant performance measure rather than a technology or infrastructure metric. Drilldown is generally implemented visually through the ability to click down through ‘layers’ of dashboards to focus on finer-grained detail. This is a valuable feature that can dramatically reduce the time and cost to diagnose incidents, specifically by allowing the relationship between a business activity and technology performance to be viewed in real-time from a single monitoring application. Drilldown can also be considered as a means to navigate the Service Model to determine where an activity is becoming degraded. The following example illustrates a practical and valuable use of this capability.

Figure 3 below outlines the architecture and flow for a common and critical business activity, namely provisioning of service to new users. For the purposes of this example, the actual service being provisioned is irrelevant; it could be a postpaid mobile user, a new insurance premium or utility service such as gas or electricity. The provisioning process starts with a request from a Point of Sales (PoS) application which may be traditional shop based, online or via a customer contact centre.

The PoS application generates a request to a provisioning gateway which forwards the request to the necessary IT systems (e.g. billing, network). The system architecture is likely to be complex and multi-layered with, for example, one or more middleware systems mediating the provisioning request to a further set of systems. In this simple example, seven ‘Units of Work’ are identified from the initial request through to the subsequent lower level middleware and application operations.

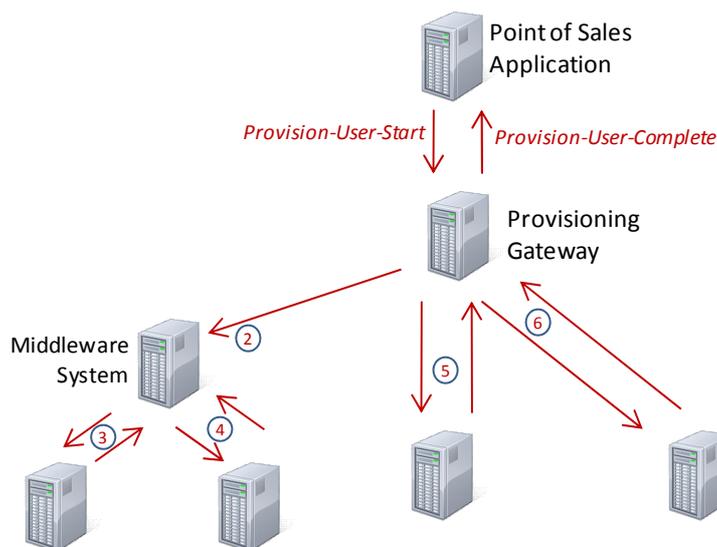


Figure 3 - Activity Definition and Units of Work

Consider the following common problem relating to this activation scenario:-

'it is taking too long to provision users, some requests are lost or timing out; this behaviour generally occurs during busy periods but also occurs at apparently random times.'

A technology driven monitoring solution would generally only alert to this situation after a failure or threshold breach in the underlying technology. Based on Figure 1 (page 3), however, it is in fact statistically improbable that the performance degradation is directly related to a simple hardware problem. If the technology-centric solution also supports monitoring the protocols or events related to the provisioning start and stop requests, then it could raise an alarm that response time targets were being exceeded. But it would provide no insight as to the root cause of the problem due to the lack of a service model that describes the set and sequence of Units of Work associated with the provisioning activity.

In contrast, the service model supported by a business-centric solution allows the provisioning activity to be defined to any required level of detail. For example, events could be monitored at each of the servers shown in Figure 3, or only a sub-set could be monitored. The decision is typically a design trade-off between the need to deploy additional sensors and the depth of insight required. Not only could the monitoring solution indicate that provisioning time was exceeding targets based on simply monitoring the Start/Complete events at the POS Application, but it could also provide insight as to where in the flow the problem was occurring. For example, there may be an interaction between the middleware system and an underlying application server that is performing poorly due to incorrect configuration or poor design. This is resulting in a bottleneck and causing the overall provisioning time to exceed target. From a technology monitoring perspective, it is highly probable that the middleware and application server appears completely healthy. **It is only with the end to end view provided by Business Activity Monitoring solutions that such insight can be achieved.**

There are two general approaches that can be used to navigate the Service Model to gain diagnostic insight:-

- **Static Mapping** - On detection of a performance issue with the business activity, the user can drill-down to inspect the IT services and technology underlying the activity. This requires a simple Service Model that maintains a mapping between the business activity and IT services. In the case of the prepaid scenario in Figure 3, the drill-down view would show key performance parameters for the POS server, middleware server, provisioning gateway and potentially even the lower level systems. For example, if the response time of the middleware system was exceeding some target, there is a high probability that the root cause of performance problems with the end-to-end business activity is related to the middleware or underlying systems.
- **Dynamic Mapping** (also known as 'Track and Trace') - In this case individual transactions are tracked through the underlying systems. Dynamic mapping requires a unique identifier to be used consistently in each Unit of Work so the business activity can be accurately correlated and tracked. This provides very powerful insight, for example allowing dynamic

reporting of the performance of each Unit of Work for every provisioning transaction and also allowing 'lost' or 'in progress' transactions to be investigated.

Conclusion

This paper has described technology-centric and business-centric approaches to monitoring for IT intensive enterprises. Analysis of the origin of incidents reveals that approximately 20% can be attributed to the type of infrastructure failure or degradation detected by technology-centric monitoring. Deployment of business-centric monitoring complements existing technology centric solutions to substantially improve incident management through earlier detection, faster diagnosis and elimination of a wider range of incident types.

WestGlobal's Vantify Experience Centre is an example of a BAM based solution that supports the pattern matching and service modelling techniques described in this paper, enabling improvements of up to 80% in incident detection rates. For more information please visit www.westglobal.com.