APM End-user Monitoring Reports User Guide

Monitoring of business transactions with the CA Wily Application Performance Monitoring (APM) Solution allows organizations to see transaction performance as their end-users experience it. In addition to raw technical performance metrics, APM allows users from within the business to view results in order to make decisions for the business.

This guide is designed to help business users understand and maximize the value of the reports within APM, and to know which reports will help illuminate challenges and successes within the business. Additional information is available from the CEM documentation, available within the product.

Included Reports

This document reviews a representative sample of the reports. Other reports use the same metrics, methodologies, and concepts, so inclusion in this document would be redundant.

Document Focus

This document will focus on two main questions: "What do the numbers mean?" and "How can I get the most value from this report?" It is not designed to be an exhaustive "how-to" documentation, as that is already available for the product in both PDF format and as context-sensitive help within the product.

Sorting Data

Any column of data in reports that can be sorted has a blue heading. Simply click the heading and the data will be sorted in ascending order. Click the heading another time, and it will be sorted in descending order.

Report Time Frames

Most of the statistical reports are produced with the default time frame of "Today." The time frame of the report can be changed to any of the following, depending on the Data Retention Settings.

- Today
- Yesterday
- Custom day (any day in the last two years of data)
- Current week
- Previous week
- Custom week (any week in the last two years of data)
- Current month
- Previous month
- Custom month (any month in the last two years of data)
- Current year
- Previous year
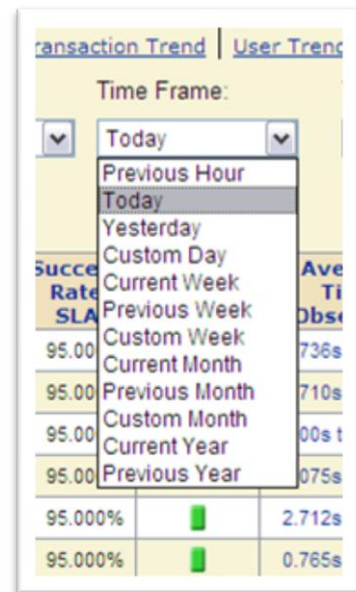
Glossary and Documentation

Rather than reprint the glossary and documentation with this report reference, find the current documentation in the following places:

- Installed CEM Appliances (behind the Wily CEM Help link):
  http://[ipaddress]/wily/cem/tess/help/administrationHelp.htm
- Public Document Site:
  https://community.wilytech.com/kbcategory.jspa?categoryID=659
- Internal (CA Wily) Document Site:
  https://wilynet.wilytech.com/entry.jspa?externalID=2926&categoryID=396

Cross-Reference: Which Reports Do I Need?

Rather than review each individual report, it may be helpful to review the problems and use cases for the reports to determine which report is the most appropriate and valuable. Review the Use cases listed below. The report and page number of that report are listed in parentheses immediately after the use case.

- Which transactions are experiencing problems? And of what type? (Incident Management Report, page 5)
- How many people are being affected by this problem? (Incident Management Report, page 5)
- Is the problem serious enough to take the application offline and fix it, or is it sufficient to fix it during the normal course of business? (Incident Management Report, page 5)
- Do we have all of the information necessary to troubleshoot and manage this incident? (Incident Detail Report, page 9)
- How long has this incident been active, and how much impact is it having? (Incident Detail Report, page 9)
- We need to reduce the time it takes to determine who needs to deal with this problem. Who should really deal with this? (Incident Troubleshooting Report, page 11)
- What is really to blame for the performance drop? Can I eliminate the need for a conference call to isolate that? (Incident Troubleshooting Report, page 11)
- What is the commonality among all of the defects that we have found? Are there specific web servers or application servers or network segments that are in many of the defects? Is there an infrastructure problem that is causing the defects? (Defect List Report, page 15)
- Which of all of the defects are the worst, and who are the users that were affected by those defects? (Defect List Report, page 15)
- What exactly happened in the defect? I need to be able to deliver the details to the triage team to help them determine the root cause of the problem. (Defect Detail Report, page 17)
- Was the defect caused by systematic issues or was it due to content issues? (Defect Detail Report, page 17)
- What time is associated with each of the layers through which the transaction flows? (Defect Detail Report, page 17)
- If the same user is experiencing the defects, are there specific details of the transaction that can help isolate the problem that are unique to this user? (Defect Detail Report, page 17)
- To ensure that our highest value customers are being well served, we need to know if any of them have been affected by our transaction incident. When we see that they have been impacted, we can call them and explain the circumstances, so we can diffuse the situation. (Affected Users Report, page 21)
- We reflect our internal SLAs in our monitoring reports. How are we doing in relation to our SLAs? (Transaction SLA Report, page 23; User SLA Report, page 27; Correlational SLA Report, page 31)

- We group our customers, and set an SLA for each. We monitor to see if we are meeting our agreed upon SLAs. Are we meeting those SLAs? What is the impact to our business if we continue to fail to meet those SLAs? (User SLA Report, page 27; Correlational SLA Report, page 31)
- We have just started a new marketing initiative. How is it reflected in the transaction numbers? How is my system performing relative to the new volume? (Transaction Trend Report, page 33)
- We have a new version of the software/database/hardware/etc. Compared to the performance before we made the change, how are we doing? (Transaction Trend Report, page 33)
- We have made a significant investment in our bug-tracking and development teams. Is the investment showing in the performance? (Transaction Trend Report, page 33)
- One of our transactions has a much lower success rate than the others. What are the factors affecting that success rate which is independent of other monitored transactions? (Performance Report, page 35; Quality Report, page 39; Count Analysis Graph, page 47)
- The average throughput for some transactions is lower than for others. Why is there a discrepancy? (Performance Report, page 35; Quality Report, page 39; Count Analysis Graph, page 47)
- How can I tell if my development process for my applications is effective? (Quality Report, page 39)
- Which transactions are having problems? (Quality Report, page 39)
- Which of our customers are at risk? How much risk is there that they might find another source? (Time Analysis Graph, page 45)
- We need to reduce the dependence our users have on the live customer support. They calling too often for us to reduce our staff. What level of confidence can we expect now? (Time Analysis Graph, page 45)
- We just added a new page to our web site, with new options for our customers. Are our customers using the new page, transactions, and capabilities? Are we having errors with the new transactions? Can people use them effectively? (Real Time Transaction Dashboards, page 49)
- We need to see how our web site and the transactions are doing at any time. We need to have a "pulse" on the counts, the transaction times, and the errors, because we need to react quickly if there is a problem. (Real Time Transaction Dashboards, page 49)

Incident Management Report

Report Location

CEM | Incident Management | Incidents

Who should view this report?

Individuals who are involved in the following areas should refer to the Performance report:

- Help desk
- Problem triage
- Problem diagnosis
- Customer service

The Problem Resolution Process

The Incident Management Report focuses on determining the cause of significant problems with transactions, providing details for diagnostics of the root cause, and recognizing who is affected by the monitored problems.
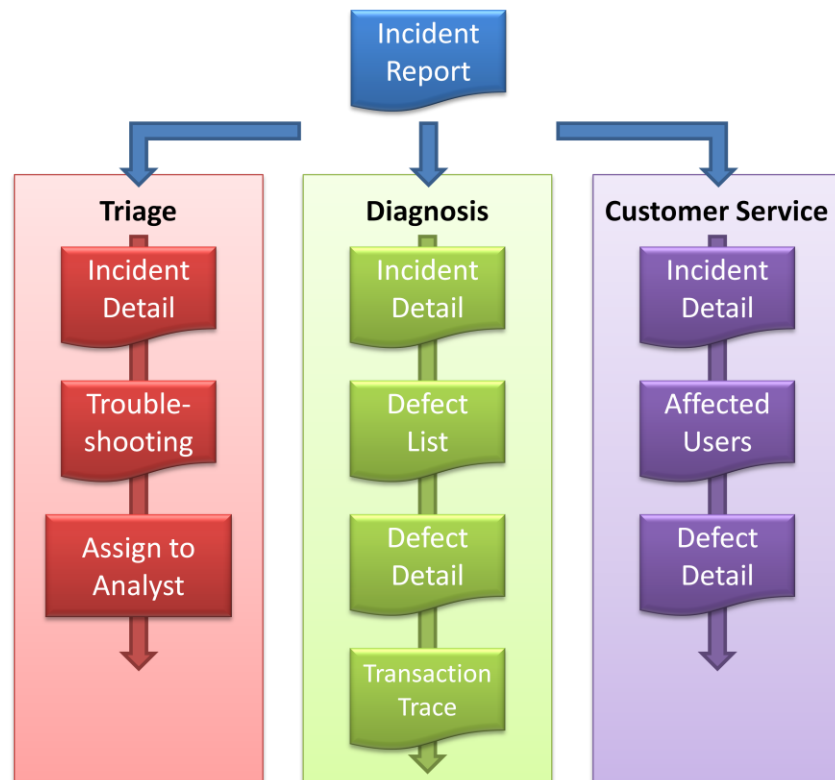
In triage, the primary goal is to determine quickly to whom the problem should be assigned. The Performance Report focuses on the actual performance of business transactions in the following areas. It can be produced at either the business process level or the business transaction level.

What do the numbers mean?

Defects are tracked based on the transaction and the type of defect. For example, login transactions with slow time defects are tracked independently of login transactions with incomplete transaction defects. Business transaction defects (termed "behavioral defects") are tracked independently of transaction component defects (also called "response defects").

Business Impact

Every defect has impact to the business, and some affect the business more than others. In order to accurately determine the impact of an incident (and its aggregated defects), each transaction, user, and type of defect must have an impact value. The values range from 0-7.

Label	Value
Ignore	0
Minimum	1
Very Low	2
Low	3
Medium	4
High	5
Critical	6
Trigger Immediately	7

Each business transaction is assigned a business impact, or it inherits the impact from the business process or global domain settings. Each user group or user is assigned an impact, or the value is inherited from the user group or the global domain settings. The impact setting for each type of defect is set at the business transaction level.

When a transaction is defective, the business impact is calculated by multiplying:

business transaction impact x **user impact** x **defect type impact**

A defective transaction with a medium transaction impact (4), a critical user (6), and a critical defect (6) would have a business impact of 144. A defective transaction with a medium transaction impact (4), a medium user (4), and a medium defect (4) would have a business impact of 64.

The business impact of the incident is the aggregated total of the business impact of all of the included defective transactions. The list of incidents in the Incident Management Report is ranked in order of business impact, indicating that the incident with the highest value has the greatest impact on the business, and therefore should be considered first.

Impact Level

The impact level is based on thresholds, and is used for escalation and notification. The default values are as follows:

Business Impact	Threshold
Low	0*
Moderate	1000
Severe	2500
Critical	5000

When the aggregate business impact of all defects in an incident reaches the moderate incident threshold, the incident shows in the Incident Management Report, and may trigger notifications to web pages, email messages to analysts, and/or communication with help desk or infrastructure monitoring systems.

*Low business impact incidents are seldom seen in the Incident Management Report. If the business impact for a business transaction, user, or type of defect is set to notify immediately, as soon as a single defect is recognized, the incident will be generated, and will appear on the list, even if the moderate impact level has not been reached.

Identified Users and Groups

Users who experience defects are tracked in each incident, and can be viewed independently or by the group to which they belong.

Incident Duration

The duration of each incident is dependent on the first recognized occurrence an instance of the defective transaction and the last recognized occurrence. Current time and the time when the incident escalated to moderate and is shown in the report have no bearing on the value.

Common Use Cases for Incident Management Report

- Which transactions are experiencing problems? And of what type?
- How many people are being affected by this problem?
- Is the problem serious enough to take the application offline and fix it, or is it sufficient to fix it during the normal course of business?

Incident Detail Report

Report Location

CEM | Incident Management | *<Incident number>*

Who should view this report?

Individuals who are involved in the following areas should refer to the Incident Detail Report:

- Help desk
- Problem triage
- Problem diagnosis
- Customer service

What do the numbers mean?

In addition to the information regarding the incident on the Incident Management Report, the Incident Detail Report includes information to help triage and manage the incident.

Included in the detail is the incident trigger time—when the incident appears on the Incident Management Report—in addition to the first and last recorded instances of defective transactions.

The Defect Time Distribution includes only timings of the defective transaction in the incident, not all monitored transactions of the type. If Introscope integration is activated, application server time, logic time, and backend time distributions will be displayed, based on information supplied by Introscope.

If the distribution of the Time to First Response is low, it is likely that the problem is not in the network infrastructure. If it is nearly the same as the overall transaction time, most of the time in the transaction is in the network.

If evidence collection is done at different escalation stages, it is listed, and the analyst may be able to click on the "yes" value and see the HTML version of the document. This allows the analyst to see the information at the time of the escalation, not just the current state of the incident.
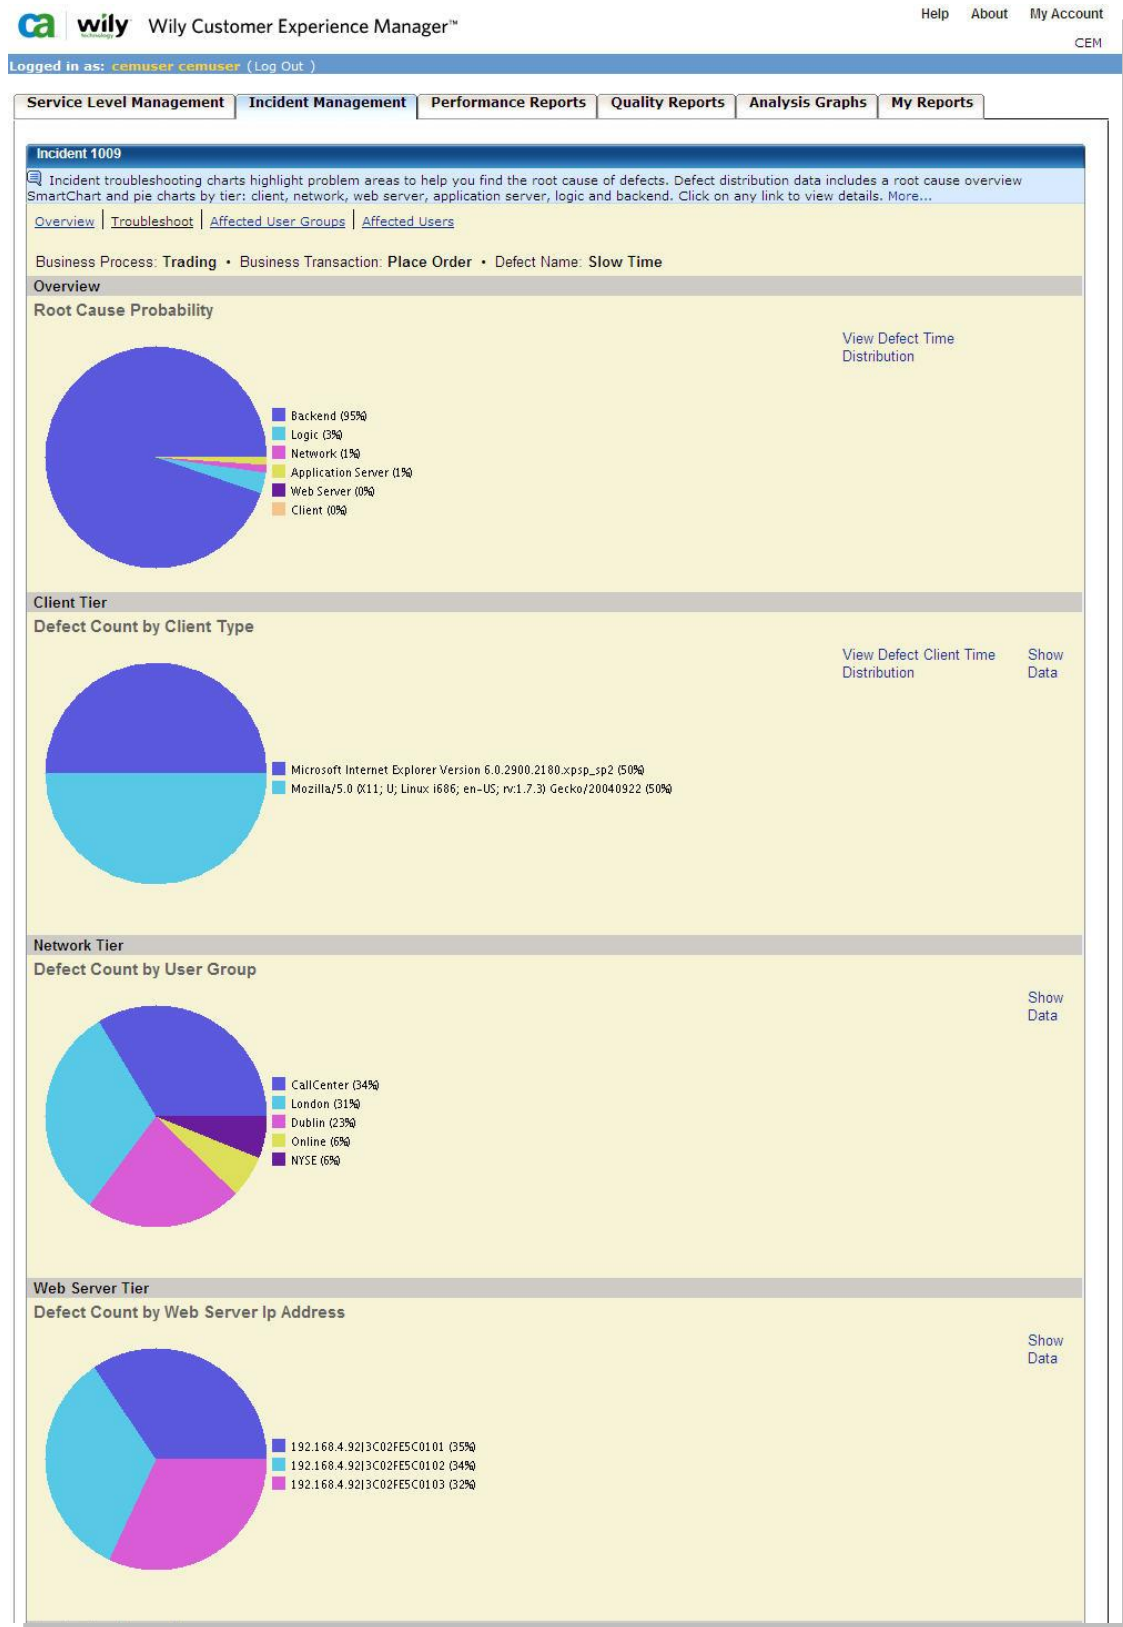
When integrated with Introscope, CEM can automatically launch transaction traces with specific details including the transaction, the duration of the trace session, and the threshold for tracing. If more traces are required, the analyst can launch them directly from within the report. If there are traces, the Incident Detail Report will show the count. Clicking on the number of traces will produce a list of defective transactions with traces, for additional triage and diagnosis. To view the traces directly, click on the links to go to Introscope.

Details on closing the incident are reflected in the Problem Resolution Cycle. This allows analysts and others the opportunity to see what was done in relation to the incident.

Common Use Cases for Incident Detail Report

- Do we have all of the information necessary to troubleshoot and manage this incident?
- How long has this incident been active, and how much impact is it having?

Incident Troubleshooting Report (Smart Chart)

Application Server Tier
Defect Count by Application Server

View Defect Application   Show
Server Time Distribution   Data

■ ForwardInc_01|App Servers|ForwardInc_01 (34%)
■ ForwardInc_03|App Servers|ForwardInc_03 (34%)
■ ForwardInc_02|App Servers|ForwardInc_02 (32%)

Logic Tier
Defect Count by Suspected Blame Component

View Defect Logic Time   Show
Distribution   Data

■ EJB|Session|TradeSessionEJB_n7enxc_impl|searchCustomerByLastNameWild (34%)
■ EJB|Session|TradeSessionEJB_n7enxc_impl|searchCustomerByUserID (34%)
■ EJB|Session|TradeSessionEJB_n7enxc_impl|searchCustomerByCategory (32%)

Backend Tier
Defect Count by Suspected Blame Backend Component

View Defect Backend   Show
Time Distribution   Data

■ WebServices|Client|http_//CustomerInfo.customer.wily.com|findCustomerByLastNameWild (100%)

Report Location

CEM | Incident Management | *<Incident number>* | Troubleshooting

Who should view this report?

Individuals who are involved in the following areas should refer to the Incident Troubleshooting Report:

- Help desk
- Problem triage
- Problem diagnosis
- Customer service

What do the numbers mean?

To help determine root cause probability, the Troubleshooting section of the Incident Management Report examines the following areas where the problem could be:

- Client tier
- Network tier (user group)
- Web server tier
- Application server tier*
- Logic tier*
- Backend tier*

*These tiers are only available when Introscope integration is activated.

The applied algorithms first eliminate the tiers where the defects occur across all possibilities, with the rationale being that if all of the web servers, for example, are equally represented in the defects, the web server is likely not the problem. The report on the previous page shows that example. However, even though the distribution may not be equal, as is shown in the Network Tier above, the distribution is among all groups. If all of the groups were represented in the transactions, but only the London group was found in the defects, it would be likely that the problem was in the network in London.

The less the distribution is equal across all possibilities, the more likely that the problem is there. Even though there were multiple web services calls made, the only one that showed up in the defects was the findCustomerByLastNameWild call.

Consequently, it is the most likely of all possibilities to be the root cause of the problems. The analyst can confidently route the incident to the correct people to have the problem diagnosed further and a solution be developed.

Common Use Cases for Incident Troubleshooting Report

- We need to reduce the time it takes to determine who needs to deal with this problem. Who should really deal with this?
- What is really to blame for the performance drop? Can I eliminate the need for a conference call to isolate that?

Defect List Report

Report Location

CEM | Incident Management | *<Incident number>* | Defects

Who should view this report?

Individuals who are involved in the following areas should refer to the Defect List Report:

- Help desk
- Problem triage
- Problem diagnosis
- Customer service

What do the numbers mean?

In every incident, there are defects. The Defect List Report shows all of the defects in the incident, sorted by the time of the defect, with the latest defect at the top of the list.

In the Customer Experience View, the client IP address, the web server IP address, and the web server MAC address are shown for each defect. Using that information, you can also determine if a network segment is experiencing problems. However, the addresses may be masked if the clients are behind a firewall, or if the web servers are behind a load balancer.

In the Introscope View, the defect information and the client addresses are replaced with information supplied by Introscope, including the application server ID, the virtual machine ID, the request handler ID, the request ID (the GUID in Introscope), and whether or not there is a transaction trace. This is particularly helpful to analysts who also have access to Introscope to follow the transaction defect information through into the Introscope console or investigator tree.

Common Use Cases for Defect List Report

- What is the commonality among all of the defects that we have found? Are there specific web servers or application servers or network segments that are in many of the defects? Is there an infrastructure problem that is causing the defects?
- Which of all of the defects are the worst, and who are the users that were affected by those defects?

Defect Detail Report

Wily Customer Experience Manager™

Logged in as: cemuser cemuser (Log Out)

Service Level Management	Incident Management	Performance Reports	Quality Reports	Analysis Graphs	My Reports

Defect: Slow Time - 23-Feb-2009 16:59:27

Details about a specific defect, where there are insufficient defects to constitute an incident, or where the defect is too recent to be included in an incident. More...

Defect Information
Defect Name:	Slow Time
Defect Condition:	8.000s
Defect Value:	10.422s
Time of Occurence:	23-Feb-2009 16:59:27
Business Impact:	112

Transaction Information
Business Transaction:	Place Order
Transaction:	
Component:	
Size:	4.9kB
Time:	10.422s
Throughput:	1.5kB/s
Time to First Response:	4.073s

User Information
User:	Blumfield, Rose
User Group:	CallCenter
Client IP Address:	192.168.20.234

Web Server Information
Web Server IP Address:	192.168.4.92
Web Server Port Number:	443
Web Server MAC Address:	3C:02:FE:5C:01:01

Application Server Information
Request ID:	11479525945RB
Application Server ID:	192.168.4.45
Virtual Machine ID:	ForwardInc_01
Request Handler ID:	ActionServlet
Application Server Hostname:	ForwardInc_01
Application Server Type:	App Servers
Application Server Time:	6.042s
Suspected Blame Component:	EJB|Session|TradeSessionEJB_n7enxc_Impl
Suspected Blame Component Time:	5.500s
Suspected Blame Backend Component:	WebServices|Client|http_//CustomerInfo.customer.wily.com|findCustomerByLastNameWild
Suspected Blame Backend Component Time:	5.132s
Detail:	Go to Introscope WebView
Detail:	Go to Introscope Workstation Web Start

Transaction Trace
Name	Duration
Customer Experience|Business Process|Trading|Business Transaction|Place Order	6.042s
Frontends|Apps|Trading|URLs|Default	6.042s
Servlets|ActionServlet	6.042s
EJB|Session|TradeSessionEJB_n7enxc_Impl	5.632s
EJB|Session|TradeSessionEJB_n7enxc_impl|searchCustomers	5.632s
EJB|Session|TradeSessionEJB_n7enxc_Impl|searchCustomersByLastNameWild	5.632s
WebServices|Client|http_//CustomerInfo.customer.wily.com	5.132s
WebServices|Client|http_//CustomerInfo.customer.wily.com|findCustomerByLastNameWild	5.132s

		HTTP Information
URL	Port:	443
URL	Path:	/Trading/searchresults.do
RequestHeader	Host:	192.168.4.45:443
URL	Host:	192.168.4.45
RequestHeader	User-Agent:	Microsoft Internet Explorer Version 6.0.2900.2180.xpsp_sp2
RequestHeader	Accept:	text/xml,application/xml,application/xhtml+xml,text/html;q=0.9,text/plain;q=0.8,image/png,*/*;q=0.5
RequestHeader	Accept-Language:	en-us,en;q=0.5
RequestHeader	Accept-Encoding:	gzip,deflate
RequestHeader	Accept-Charset:	ISO-8859-1,utf-8;q=0.7,*;q=0.7
RequestHeader	Keep-Alive:	300
RequestHeader	Connection:	keep-alive
RequestHeader	Referer:	http://192.168.4.45:443/Trading/searchresults.do
Cookie	Language:	en
Cookie	Country:	US
Cookie	JSESSIONID:	G4XLthkJFyyHm9BJmJ1FdvKIxnCJyLQzWrkp71KIXNFfp8hrH2nb!400988984
Cookie	wily_charEncoding:	UTF-8
Cookie	wily_clickTime:	1178126790780
Cookie	wily_clickLink:	http%3A%2F%2F192.168.1.202%3A7000%2FTrading%2Fsearchresults.do
RequestHeader	Cache-Control:	max-age=0
ResponseHeader	Date:	23-Feb-2009 17::01 GMT
ResponseHeader	Pragma:	No-cache
ResponseHeader	Content-Type:	text/html;charset=UTF-8
ResponseHeader	Expires:	Thu, 01 Jan 1970 00:00:00 GMT
ResponseHeader	x-wily-servlet:	Encrypt1 WSry9LDfbZEeaCzdIYhs5YQASIgB5TXCASrogQVGID0Wpu+6bdm1JeVOSnMKI5uCTmW7c1QiCABcpWb39vm+JpzPCWnczvW
ResponseHeader	x-wily-info:	Clear guid=4E10809DC0A801CA00B15D02350018A2
ResponseHeader	Transfer-Encoding:	chunked
ResponseHeader	Cache-Control:	no-cache
ResponseBody	body:	Click here
RequestHeader	Content-Length:	305
Post	lastName:	Blumfield
Post	firstName:	Rose
Post	email:	Rose.Blumfield@wily.com
Post	Cust ID:	******
Post	Pin:	******
Post	action:	Place Order
ResponseHeader	Connection:	close

Report Location

CEM | Incident Management | *<Incident number>* | Defects | *<Defect time stamp>*

Who should view this report?

Individuals who are involved in the following areas should refer to the Performance report:

- Help desk
- Problem triage
- Problem diagnosis
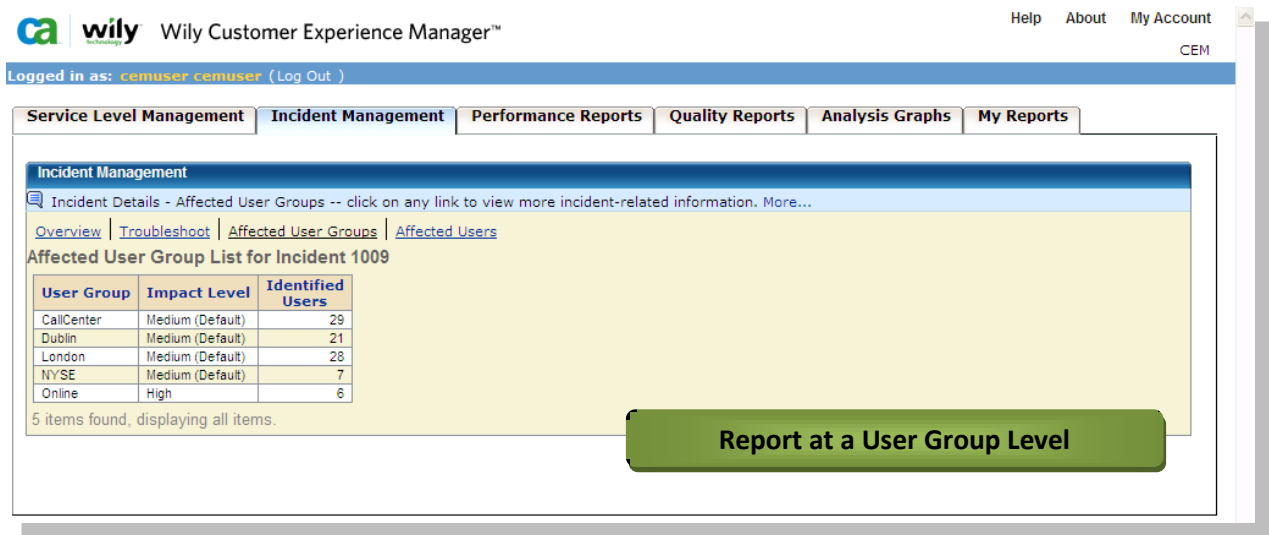- Customer service

What do the numbers mean?

The HTTP information (shown on the previous page) is optional and dependent on the license used. The HTTP-enabled license will allow the analyst to see all of the information captured during the transaction. The standard license does not allow the analyst to see any of the HTTP request/response body. Changing the license (a no-cost option) is required to change the view.

Common Use Cases for Defect Detail Report

- What exactly happened in the defect? I need to be able to deliver the details to the triage team to help them determine the root cause of the problem.
- Was the defect caused by systematic issues or was it due to content issues?
- What time is associated with each of the layers through which the transaction flows?
- If the same user is experiencing the defects, are there specific details of the transaction that can help isolate the problem that are unique to this user?

Affected User Groups/Affected Users Report

Report Location

CEM | Incident Management | Affected User Groups/Affected Users

Who should view this report?

Individuals who are involved in the following areas should refer to the Affected Users Report:

- Help desk
- Problem triage
- Problem diagnosis
- Customer service

What do the numbers mean?

Users are identified by CEM using their user ID. All users are automatically assigned to the New Users group, unless the system has been configured to assign them based on information that is passed during the transaction or based on LDAP grouping.

If the groups are assigned geographically, the Affected User Group Report can be a helpful way to quickly determine if the problem is related to a specific geographic network.

The Affected Users Report can be helpful for the customer service representatives, or even executives, so they can contact the high-value users and explain what is happening prior to the customer contacting support. The report is automatically sorted first by impact level and then alphabetically by login name.

Logged in as: cemuser cemuser (Log Out)

| Service Level Management | Incident Management | Performance Reports | Quality Reports | Analysis Graphs | My Reports |

Incident Management

Incident Details - Affected Users -- click on any link to view more incident-related information. More...

Overview | Troubleshoot | Affected User Groups | Affected Users

Affected User List for Incident 1009

User	Login Name	Impact Level	User Group
Thorson, Jane	jthorson	Critical	Online
Blumfield, Rose	rblumfield	Critical	CallCenter
Green, Silas	sgreen	High	Online
Jellico, Mikeal	mjellico	High	Online
McIlroy, Dermot	dmcilroy	High	Online
Tuomo, Jesper	jtuomo	High	Online
Alba, Fiona	falba	Medium (Default)	Dublin
Allgood, Stephanie	sallgood	Medium (Default)	CallCenter
Arispe, Marina	marispe	Medium (Default)	CallCenter
Bean, Jason	jbean	Medium (Default)	CallCenter
Berger, Eric	eberger	Medium (Default)	CallCenter
Blumfield, Rose	rblumfield	Medium (Default)	CallCenter
Braun, Matt	mbraun	Medium (Default)	London
Bundy, Curtis	cbundy	Medium (Default)	CallCenter
Christian, Val	vchristian	Medium (Default)	CallCenter
Cohn, Marty	mcohn	Medium (Default)	Dublin
Cox, Simeon	scox	Medium (Default)	London
Ellicot, Katherine	kellicot	Medium (Default)	London
Emi, Kiyoshi	kemi	Medium (Default)	NYSE
Emoy, Jasmine	jemoy	Medium (Default)	Dublin
Endoh, Mayuri	mendoh	Medium (Default)	Dublin
Ermine, Art	aermine	Medium (Default)	Dublin
Esmay, Terrence	tesmay	Medium (Default)	Dublin
Fujino, Kiichi	kfujino	Medium (Default)	NYSE
Gentry, Harrison	hgentry	Medium (Default)	London

91 items found, displaying 1 to 25.   [First/Prev] 1, 2, 3, 4 [Next/Last]

Report at an Individual User Level

Common Use Cases for Affected User Groups/Affected Users Report

- To ensure that our highest value customers are being well served, we need to know if any of them have been affected by our transaction incident. When we see that they have been impacted, we can call them and explain the circumstances, so we can diffuse the situation.

Transaction SLA Report

Report Location

CEM | Service Level Management | Transaction SLA

Who should view this report?

Individuals who are involved in the following areas should refer to the Transaction SLA Report:

- IT application owner
- Performance management
- Quality assurance

What do the numbers mean?

The Transaction SLA Report focuses on the projected service level agreements (SLAs) for business transactions in the following areas.

- Targeted Success Rate
- Average Response Time

The Success Rate is the percentage of defect-free transactions. Defects include slow transaction times, incomplete transactions, missing components and other problems, though being defective does not necessarily mean that the transaction did not complete. The Success Rate SLA is a target for the defect-free transactions. At the business process level, this is an average success rate for all transactions within the business process.

The Average Response Time is the average of all transactions in a specific business process or for a specific business transaction. This is independent of the defect threshold for business transactions. The defect threshold represents a value, where if the transaction is slower, the problem resolution team will be involved. The average time SLA represents the actual target time of the transaction, and may be significantly different than the defect threshold.

Also, there is no notification of a violation of the average time SLA. At the business process level, there may be a range of values, representing the average times of the individual business transactions within the business process.

Each business process and business transaction can either inherit the SLAs from the domain-level settings, or can be set individually.

The results are ranked based on the success rate, with the lowest success rate on top, assuming that is the area of greatest concern.

The time frame may also be adjusted, based on available time frames. (See page x for the list and explanation of possible time frames)

Common Use Cases for Transaction SLA Report

- We reflect our internal SLAs in our monitoring reports. How are we doing in relation to our SLAs?

User SLA Report

Report Location

CEM | Service Level Management | User SLA

Who should view this report?

Individuals who are involved in the following areas should refer to the User SLA report:

- IT application owner
- Business manager
- Account manager
- Performance management
- Quality assurance

What do the numbers mean?

The calculations for the User SLA report mirror the calculations for the Transaction SLA Report. However, instead of calculating the SLAs based on transactions, it is calculated by user group and individual user. An SLA set for a user or group applies to all transactions, and is independent of the transaction SLA.

For example, a transaction has an SLA of 4 seconds, and the user has a user SLA of 8 seconds. An instance of the transaction by the user executes in 6 seconds. It would fail when compared against the transaction SLA, but would succeed when measured against the user SLA.

Each user and user group can either inherit the SLAs from the domain-level settings, or can be set individually.

The results are ranked based on the success rate, with the lowest success rate on top, assuming that is the area of greatest concern.

The time frame may also be adjusted, based on available time frames. (See page x for the list and explanation of possible time frames)
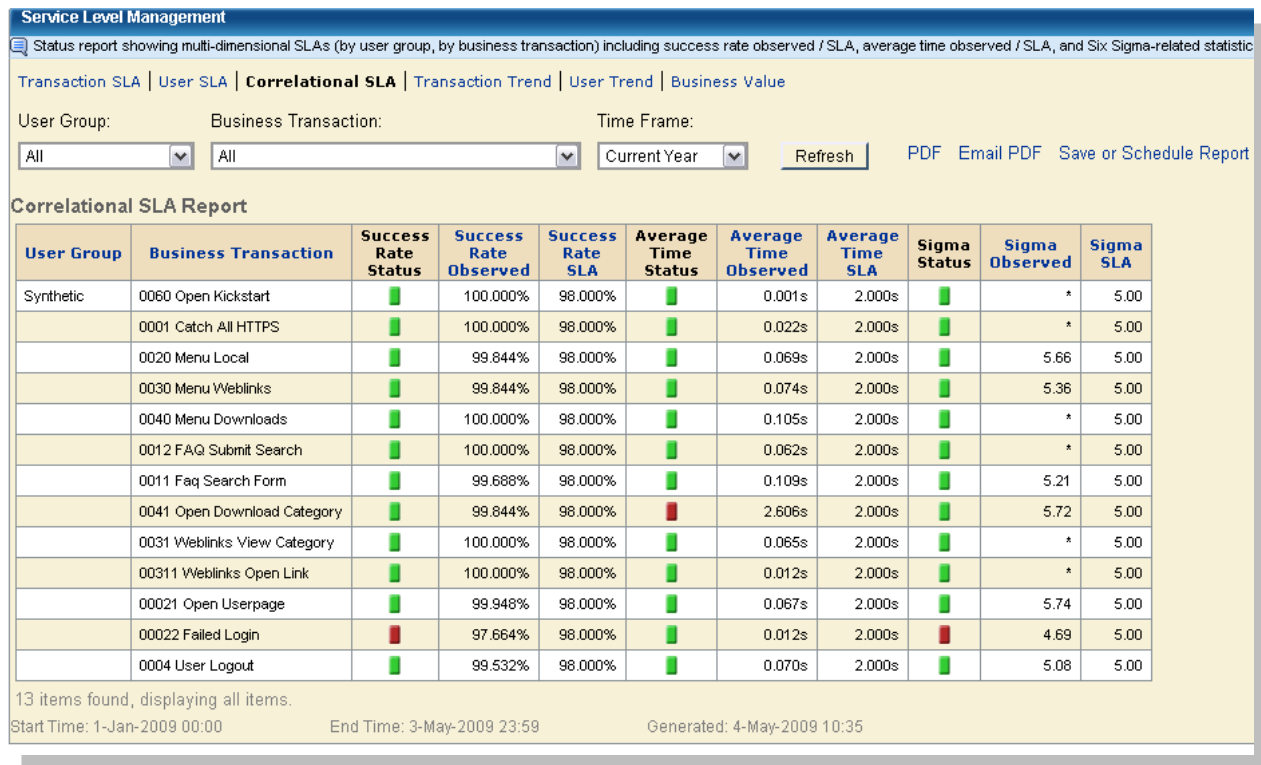
Common Use Cases for User SLA Report

- We reflect our internal SLAs in our monitoring reports. How are we doing in relation to our SLAs?

- We group our customers, and set an SLA for each. We monitor to see if we are meeting our agreed upon SLAs. Are we meeting those SLAs? What is the impact to our business if we continue to fail to meet those SLAs?

Correlational SLA Report

Service Level Management

Status report showing multi-dimensional SLAs (by user group, by business transaction) including success rate observed / SLA, average time observed / SLA, and Six Sigma-related statistic

Transaction SLA | User SLA | **Correlational SLA** | Transaction Trend | User Trend | Business Value

User Group: [All ▼]   Business Transaction: [All ▼]   Time Frame: [Current Year ▼]   [Refresh]   PDF Email PDF Save or Schedule Report

Correlational SLA Report

User Group	Business Transaction	Success Rate Status	Success Rate Observed	Success Rate SLA	Average Time Status	Average Time Observed	Average Time SLA	Sigma Status	Sigma Observed	Sigma SLA
Synthetic	0060 Open Kickstart	▌	100.000%	98.000%	▌	0.001s	2.000s	▌	*	5.00
	0001 Catch All HTTPS	▌	100.000%	98.000%	▌	0.022s	2.000s	▌	*	5.00
	0020 Menu Local	▌	99.844%	98.000%	▌	0.069s	2.000s	▌	5.66	5.00
	0030 Menu Weblinks	▌	99.844%	98.000%	▌	0.074s	2.000s	▌	5.36	5.00
	0040 Menu Downloads	▌	100.000%	98.000%	▌	0.105s	2.000s	▌	*	5.00
	0012 FAQ Submit Search	▌	100.000%	98.000%	▌	0.062s	2.000s	▌	*	5.00
	0011 Faq Search Form	▌	99.688%	98.000%	▌	0.109s	2.000s	▌	5.21	5.00
	0041 Open Download Category	▌	99.844%	98.000%	▌	2.606s	2.000s	▌	5.72	5.00
	0031 Weblinks View Category	▌	100.000%	98.000%	▌	0.065s	2.000s	▌	*	5.00
	00311 Weblinks Open Link	▌	100.000%	98.000%	▌	0.012s	2.000s	▌	*	5.00
	00021 Open Userpage	▌	99.948%	98.000%	▌	0.067s	2.000s	▌	5.74	5.00
	00022 Failed Login	▌	97.664%	98.000%	▌	0.012s	2.000s	▌	4.69	5.00
	0004 User Logout	▌	99.532%	98.000%	▌	0.070s	2.000s	▌	5.08	5.00

13 items found, displaying all items.

Start Time: 1-Jan-2009 00:00   End Time: 3-May-2009 23:59   Generated: 4-May-2009 10:35

Report Location

CEM | Service Level Management | Correlational SLA

Who should view this report?

Individuals who are involved in the following areas should refer to the Correlational SLA report:

- Performance management
- Quality assurance

What do the numbers mean?

There are SLAs for transactions and SLAs for users and groups. Normally, they are independent of each other, but Correlational SLAs allow you to create SLAs that include both components. Each group or user can have a specific SLA for each transaction, allowing you to tweak the reporting and monitoring to ensure customer satisfaction.

An example would be to set the SLA for the London group and the trading transactions to 6 seconds, and the portfolio transactions to 8 seconds. The SLA for the New York group would be 3 seconds for the trading transactions and 4 seconds for the portfolio transactions.

The time frame may also be adjusted, based on available time frames. (See page x for the list and explanation of possible time frames)

Common Use Cases for Correlational SLA Report

- We reflect our internal SLAs in our monitoring reports. How are we doing in relation to our SLAs?
- We group our customers, and set an SLA for each. We monitor to see if we are meeting our agreed upon SLAs. Are we meeting those SLAs? What is the impact to our business if we continue to fail to meet those SLAs?

Transaction Trend Report

Report Location

CEM | Service Level Management | Transaction Trend

Who should view this report?

Individuals who are involved in the following areas should refer to the Transaction Trend Report:

- Resource allocation
- Capacity planning
- Quality assurance
- Performance management

What do the numbers mean?

The Transaction Trend Report focuses on the change from one time period to another, showing improvement/deterioration of Success Rate, average time, and transaction volume.

- If the trend is worse (lower success rate or higher average time), the arrows will point down and will be red ⬇.
- If the trend is better (higher success rate or lower average time), the arrows will point up and will be green ⬆.
- A higher or lower count is not necessarily good or bad, so the blue arrow ⬇ reflects the direction of the trend, but does not reflect good or bad.
- If there is no change, no arrow is shown.

The key to this report is to compare either the Success Rate Trend or the Average Time Trend to the Transaction Count Trend.

- If the success rate or the average time is worse, and the transaction count is higher, it may be a capacity problem, and resources should be allocated to capacity planning.
- If the success rate or the average time is worse, and the transaction count is the same or lower, it is probably *not* a capacity problem, but is likely an application or network problem. Resources should be allocated to determine where the problem is (see Incident Management Report).
- If the success rate or the average time is worse or unchanged, and the transaction count is higher, there is sufficient capacity to handle the load, and prior planning for the extra load has been effective.
- If the transaction count is significantly lower, there may be a problem with the transaction's page, there may be an external reason why fewer people are using the transaction, or the transaction may have changed so the monitoring needs to be adjusted.

The "Primary Time Frame" is reflected in the "Current" success rate, average time, and transaction count. The "Secondary Time Frame" is shown as the "Previous" rate, time, and count. Normally, the more current time will be the primary time frame, and the older time will be the secondary time frame. (See page x for the list and explanation of possible time frames)

Common Use Cases for Transaction Trend Report

- We have just started a new marketing initiative. How is it reflected in the transaction numbers? How is my system performing relative to the new volume?
- We have a new version of the software/database/hardware/etc. Compared to the performance before we made the change, how are we doing?
- We have made a significant investment in our bug-tracking and development teams. Is the investment showing in the performance?

Performance Report

Report Location

CEM | Performance Reports | Performance Report

Who should view this report?

Individuals who are involved in the following areas should refer to the Performance Report:

- Performance management
- Quality assurance

What do the numbers mean?

The Performance Report focuses on the actual performance of business transactions in the following areas. It can be produced at either the business process level or the business transaction level.

- Success Rate
- Transaction Counts
- Averages, based on all monitored transactions, including
  - Throughput
  - Time to complete the transaction
  - Size of transactions, in kilobytes
- Total volume of all monitored transactions

The success rate is the percentage of defect-free transactions. Defects include slow transaction times, incomplete transactions, missing components, and other problems, though being defective does not necessarily mean that the transaction did not complete. The items in the report are ranked by the success rate, with the item with lowest rate of success at the top. At the business process level, this is an average success rate for all transactions within the business process.

The transaction counts show the number of transactions monitored, and segments them into good and defective transactions. Good transactions have no recognized defects.

Throughput is the measure of flow of data, with higher numbers being better than lower. Each transaction is measured for throughput, and the average throughput displayed is the average of all of the throughputs from each transaction, rather than the calculation of the value based on the average time and average size displayed in the report. Throughput may vary depending on the systems producing the data. For example, since calculation time is also included in this value, a transaction that includes web services or includes complex calculations may take longer than those that are simple transactions.

The average time is a simple average of the total transaction time.

The size of the transaction is the total size from the first byte of the first request to the last byte of the last response. The average size is a simple average of the monitored size of each transaction. The total volume is a sum of the values of the sizes of all monitored transactions.

For similar transactions, reviewing the transaction times may be sufficient to determine if there are poorly performing transactions. However, viewing the throughput is helpful when the results of the transaction vary significantly in size. For example, a search transaction that returns a small subset of data will likely take less time than one that returns a large subset of the data. In that case, throughput is a better indicator of performance than average time.

The values in the performance report must be compared to an expectation of performance or a historical result. In the absence of historical or expected values, the best use of the data is to look for abnormalities in the results.

The time frame may also be adjusted, based on available time frames. (See page x for the list and explanation of possible time frames)

Common Use Cases for Performance Report

- One of our transactions has a much lower success rate than the others. What are the factors affecting that success rate that is independent of other monitored transactions?
- The average throughput for some transactions is lower than for others. Why is there a discrepancy?

Quality Report

Report Location

CEM | Quality Reports | Quality Report

Who should view this report?

Individuals who are involved in the following areas should refer to the Quality report:

- Performance management
- Quality assurance
- Six Sigma analysts
- Development teams

What do the numbers mean?

The Quality Report focuses on Six Sigma-based statistics.

The foundation of the statistical calculations is based on the number of potential defects in any given transaction, or the number of opportunities. For example, a transaction may have seven (or more or less) potential defects, including slow time, low throughput, missing transaction, partial transaction, missing component, server error, and client error. In this instance, there would be seven opportunities for success or failure.

The yield is the percentage of opportunities that are successfully completed. If one of the defects above (out of seven potential) was recognized, the yield for that transaction would be 85.7%. If two defects were detected for the same instance of the transaction, the yield would be 71.4%. Additional defects within an individual transaction reduce the yield for that transaction. The quality report is sorted by yield, with the lowest yield at the top of the list.

This differs from the success rate shown in the performance report. Regardless of the number of defects in a single transaction, the entire transaction is defective. In the performance report, the transaction is either defective or it isn't. The quality report shows how defective the transaction is.

DPMO is an acronym for defects per million opportunities. Some processes have many opportunities; others have very few. DPMO normalizes the value, so it is possible to compare disparate processes to determine which should be addressed first. The process with the highest DPMO is the one that has the lowest quality.

The Sigma value represents the calculation of standard deviations from the mean. APM uses the same 1.5 Sigma shift as standardized by Six Sigma. The table below shows the number of DPMO and the yield at several Sigma levels.

Sigma Level (σ)	DPMO	Yield
Infinite	0	100%
6	3.4 or less	99.9999%
5	230 or less	99.98%
4	6200 or less	99.38%
3	67,000 or less	93.30%

In contrast to the Performance Report and its averages, the Quality Report calculates the median values for the monitored transactions.

The median is the middle value of all recognized values, with 50% of the results lower, and 50% of the results higher. Extremely high values can skew the average, whereas the median essentially ignores the value of the outliers in the distribution. Consequently, the median can be a better representative of the middle of the performance in applications, since there are only high value outliers without low value outliers to balance the calculation.

Report at a Business Transaction Level

The time frame may also be adjusted, based on available time frames. (See page x for the list and explanation of possible time frames)

Common Use Cases for Quality Report

- How can I tell if my development process for my applications is effective?
- Which transactions are having problems?

Pareto Analysis Graph
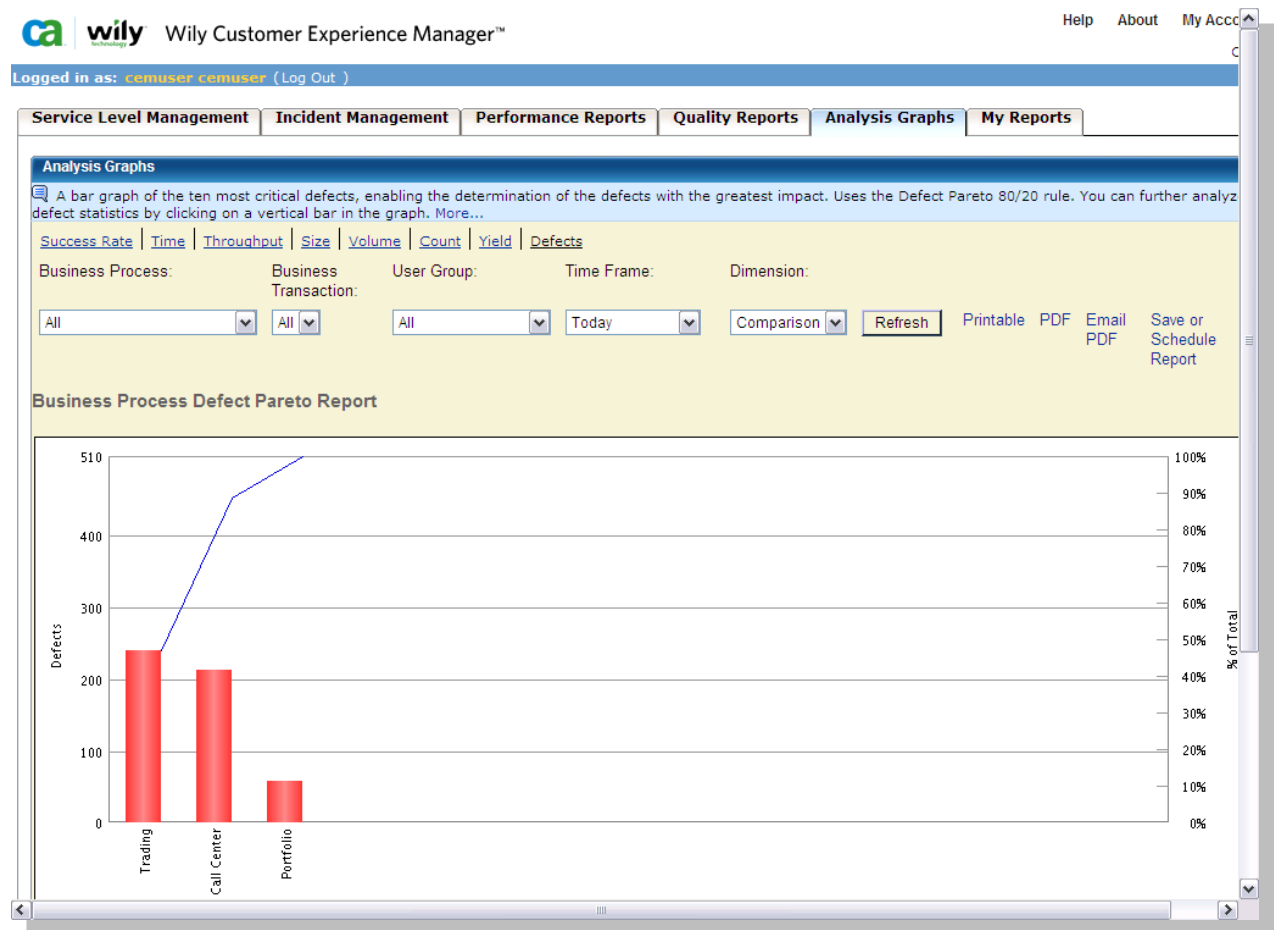
Report Location

CEM | Analysis Graphs | Defects

Who should view this report?

Individuals who are involved in the following areas should refer to the Pareto Analysis Graph:

- Help desk
- Problem triage
- Problem diagnosis
- Customer service

What do the numbers mean?

The Pareto is a bar chart that is used to add context to the information in the other reports. It is based on the Pareto Principle which states that 20% of the problems have 80% of the impact. Consequently, it is critical to know which problems have the greatest impact, not just which ones occur most often.

The Pareto helps visually recognize which problems are the most important. It gets its data from the business impact, rather than the defect count. Depending on the impact calculated for each defect, other transactions may have a larger count, but a lower impact.

The blue line bends at the mark representing 80% of the impact. The farther left the bend is, the more concentrated the problems are in a few transactions. The farther to the right the bend is, the more general and common the problems are.

Common Use Cases for Pareto Analysis Graph

- One of our transactions has a much lower success rate than the others. What are the factors affecting that success rate that is independent of other monitored transactions?
- The average throughput for some transactions is lower than for others. Why is there a discrepancy?

Time Analysis Graph
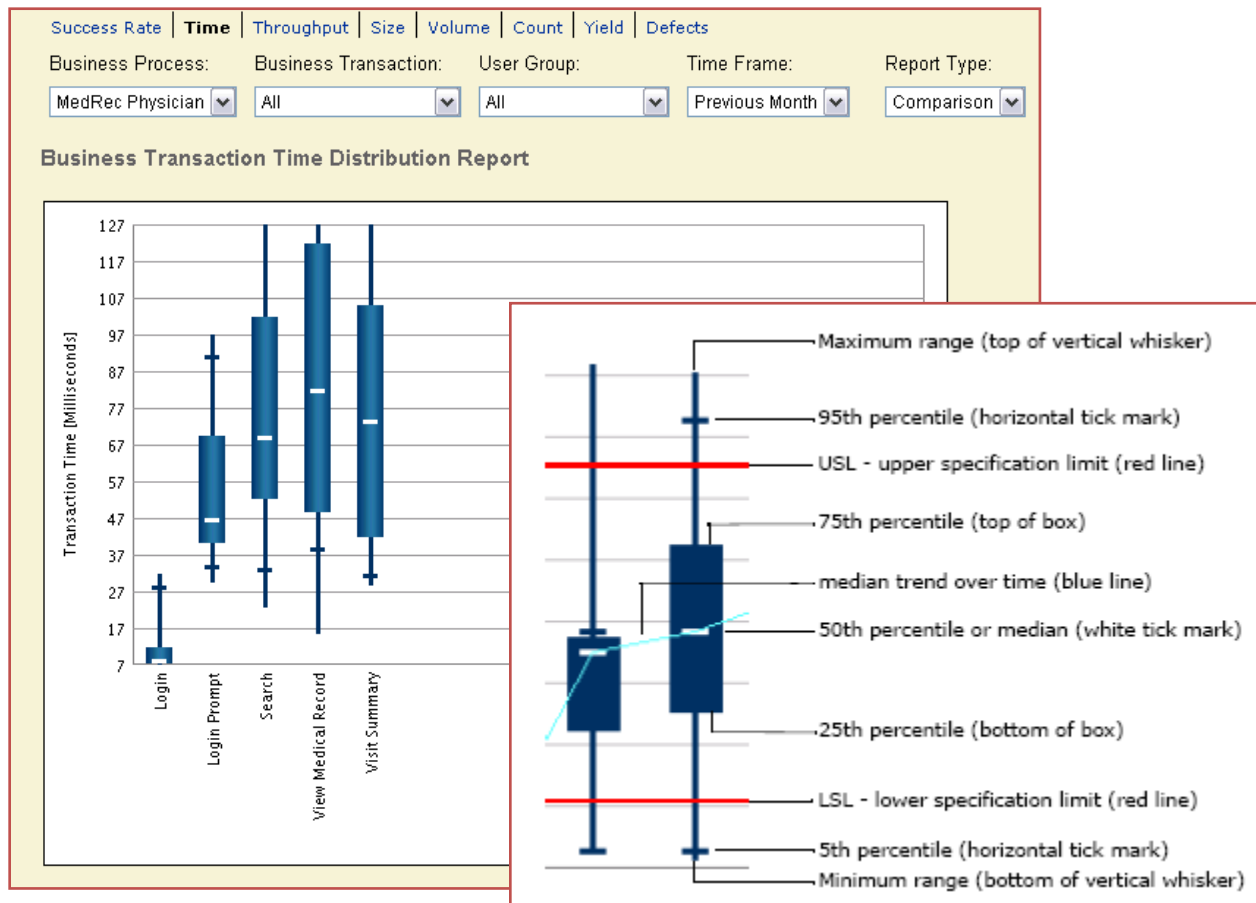
Report Location

CEM | Analysis Graphs | Time

Who should view this report?

Individuals who are involved in the following areas should refer to the Time Analysis Graph:

- Performance management
- Quality assurance
- Risk assessment

What do the numbers mean?

The Time Analysis Graph is a box whisker chart, representing a comparison of distribution curves, as if you were looking down from above.

- The top and bottom represent the minimum and maximum values recognized
- The tick marks represent the 5th and 95th percentiles
    - 5% of all transaction times were below the lower tick mark
    - 5% of all transaction times were above the upper tick mark
- The box represents the 25th-75th percentiles of transaction times
- The white line in the box represents the median value recognized

The values above the 95th percentile indicate risk. The users that experienced that 5% of the transactions are at risk to choose another vendor or method for their transaction, resulting in lost revenue or increased costs. The higher values in that 5%, the greater the risk.

The blue box represents the majority of the users and their experience. If there is a large box, representing a large variance in the performance of the transactions, it is likely that there will be a lack of confidence in the performance of the system. If one time a transaction takes 3 seconds, but another time it takes 30 seconds, and another time it takes 45 seconds, and another time it takes 5 seconds, the user will not know if the transaction completed, didn't complete, or even if it will. That erosion of confidence will prompt the user to select a different method of executing the transaction, or a different vendor.

If the web-based transactions are replacing brick-and-mortar or other high-cost/labor-intensive systems, it may be impossible to reduce those services and their attendant costs.

Common Use Cases for Time Analysis Graph

- Which of our customers are at risk? How much risk is there that they might find another source?
- We need to reduce the dependence our users have on the live customer support. They calling too often for us to reduce our staff. What level of confidence can we expect now?

Count Analysis Graph
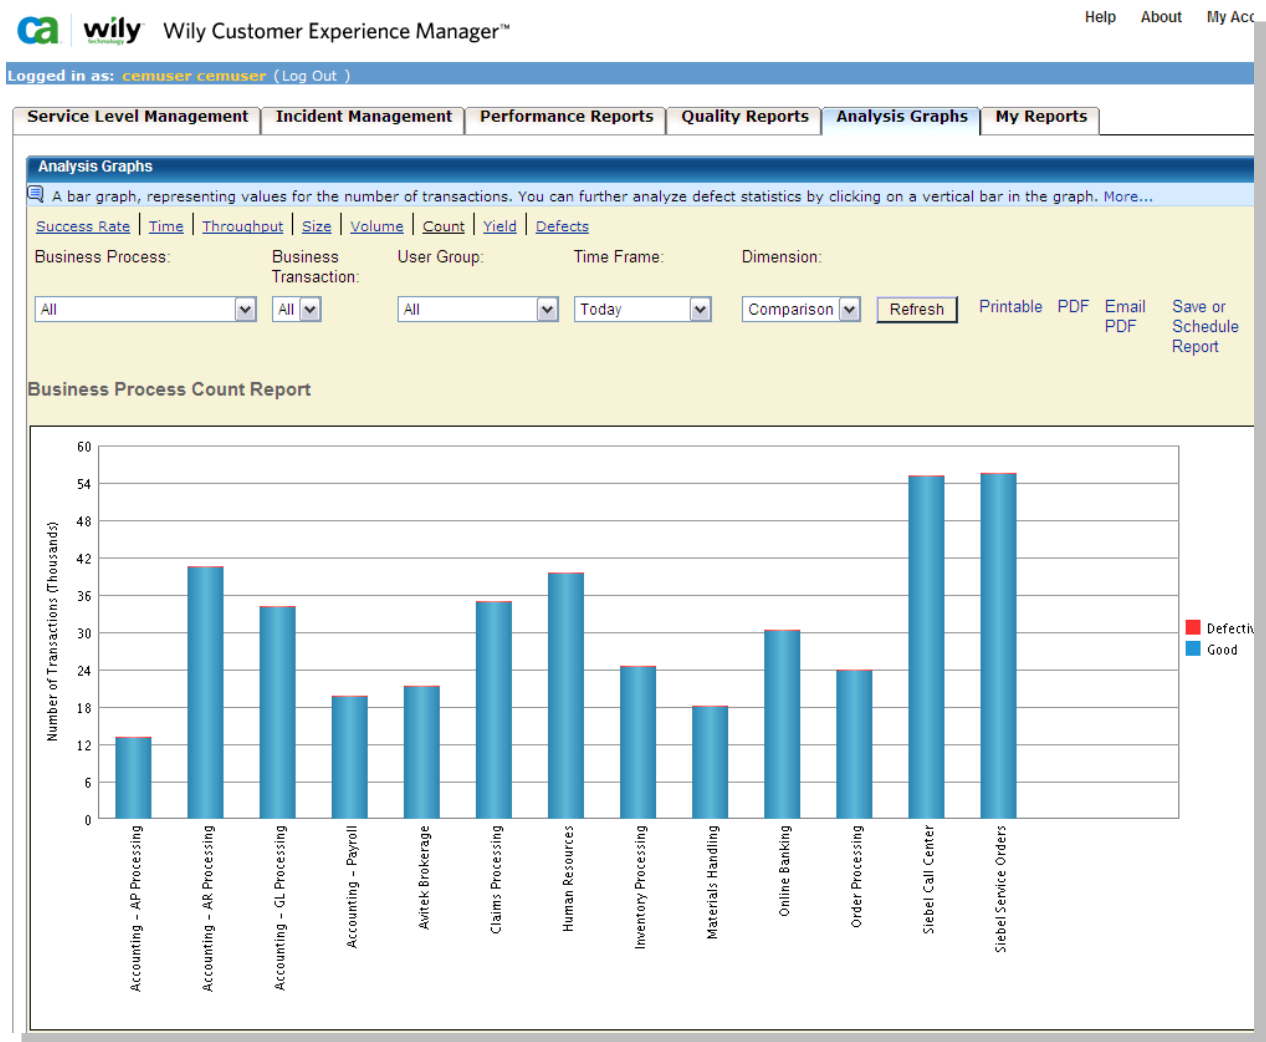
Report Location

CEM | Analysis Graphs | Count

Who should view this report?

Individuals who are involved in the following areas should refer to the Count Analysis Graph:

- Help desk
- Problem triage
- Problem diagnosis
- Customer service

What do the numbers mean?

The Count Chart is a stacked bar chart that shows the counts of the transactions. The blue portion represents the defect-free transactions and the red portion represents the defective transaction. It is easy to see how the number of transactions vary and whether the percentage of defective transactions is significant.

Common Use Cases for Count Analysis Graph

- One of our transactions has a much lower success rate than the others. What are the factors affecting that success rate that is independent of other monitored transactions?
- The average throughput for some transactions is lower than for others. Why is there a discrepancy?

Real Time Transaction Metrics
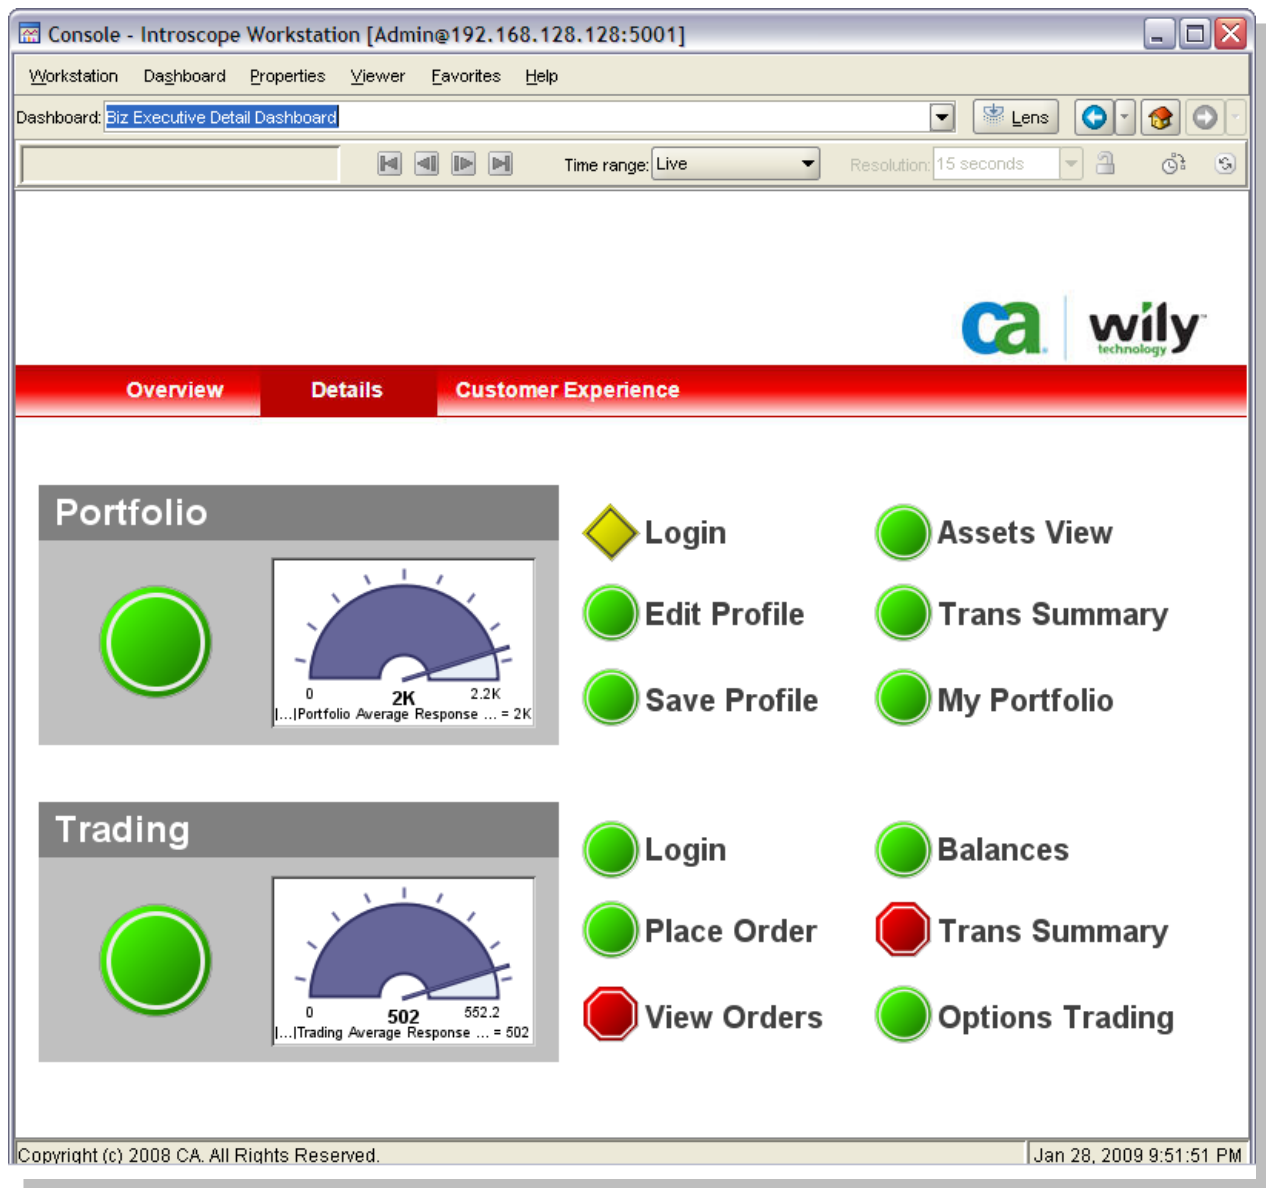
Report Location

Introscope Console

Who should view this report?

Individuals who are involved in the following areas should refer to the Performance report:

- Executive management who need to see the results of recent initiatives, until the new item has been confirmed to be stable
- Problem triage
- Problem diagnosis
- Customer service

What do the numbers mean?

Real Time Transaction Metrics shows CEM data in an Introscope console. However, rather than the typical hourly reports, the following data is sent to the console every 15 seconds.

- Average response time
- Transaction counts per interval
- Defects per interval

With that information, it is easy to see a "snapshot' of the performance of the system in near real-time. If there are problems, errors, or even usability issues, it is easy to see that in a very quick time frame.

The flexibility of the console allows rapid changes to be made, and the language of the dashboard can reflect the organization's needs. Links can be created to allow deeper investigation in Introscope, or context from other CEM reports.

Common Use Cases for Real Time Transaction Dashboards

- We just added a new page to our web site, with new options for our customers. Are our customers using the new page, transactions, and capabilities? Are we having errors with the new transactions? Can people use them effectively?
- We need to see how our web site and the transactions are doing at any time. We need to have a "pulse" on the counts, the transaction times, and the errors, because we need to react quickly if there is a problem.