

Tech Tips and Blogs 2014-2015

.v01d 12/16/2015

Hallett German

hallett.german@ca.com

CA Technologies

Distribution Notes: This document may be freely distributed to all interested audiences to help APM Administrators and Users. However, no modifications may be made to this work.

Table of Contents

Section 1: Introduction	2
Section 2: Blogs 2014-2015.....	2
2.1 It's the Architecture Silly (2014).....	2
2.2 Another look at APM System Administration Philosophies (2014).....	3
2.3 Four Keys to Successful APM Searches: Part One (2014)	5
2.4 Four Keys to Successful APM Searches: Part Two (2014).....	6
2.5: The Loneliness of the Solitary APM Administrator (2015).....	8
2.6 The Wonder of APM CE Services (2015).....	10
2.7 Problem Solving: Six Ways to Avoid the Wrong Path (2015)	12
2.8 APM Administrators: A Dying Breed? (2015).....	14
2.9 The Confessions of a Serial Change Agent: Part 1 Overview	15
2.10 The Confessions of a Serial Change Agent: Part 2 APM.....	17
Section 3: Tech Tips (2015).....	19
3.1 APM CE Knowledge Levels	19
3.2 Planning to Monitor your Application by APM CE	20
3.3 APM CE Technical Tip Overview.....	22
3.4 Tim Settings Part 1	23
3.5 Tim Settings Part 2	26
3.6 Tim Settings Part 3	27
3.7 E-Commerce and User Groups	29
3.9 APM Database Corruption	33
3.10 Ten Central Myths of APM Administrators	35

3.11 Successfully Configuring and Troubleshooting RTTM with User Groups	37
3.12 TIM Trace Options -- The Missing Pages Part 1	39

Section 1: Introduction

It is another year and I continue the “new tradition” of making an e-book of all of my blogs/tech tips. I release this “AS-IS “with minimal proofing. I hope this is of use to someone and wish you all a happy 2016.

Hallett German

CA Technologies APM Support

Hallett.german@ca.com

Section 2: Blogs 2014-2015

2.1 It's the Architecture Silly (2014)

URL: <https://communities.ca.com/people/gerha02/blog/2014/12/06/its-the-architecture-silly>

Introduction

In my time with APM, I have seen both sides of the coin -- extensive architectural planning for APM deployment and the business impact of failing to do so.

Here are some suggestions and common scenarios on this overlooked topic:

The Proactive Approach

- Takes the time to appropriately define APM requirements for current and forecasted changes for the forthcoming future.
- Reviews one to four times a year their architecture against present and future requirements.
- Only captures the metrics, defects, etc. that it needs. Does not go with default APM configurations.
- Ties the above to overall monitoring/APM deployment strategy.

The Minimum Administration/Reactive Approach

- Keeps adding Introscope agents, monitored business applications etc. until one day things breaks.
- Captures all metrics turned on by default whether needed or not just because they *may* be needed some day.
- Has that fateful day when performance and functionality is impacted and must review architecture reactively. Much like the unwary boiling frog.

http://en.wikipedia.org/wiki/Boiling_frog

- Repeats the above cycle
- Has no APM deployment/roadmap strategy

Troublesome Architecture Scenarios

There are many times that I see one of the following situations:

* No/minimal Architecture Planning was performed during deployment. CA Technologies may not have sanity-checked the final architecture.

* An Architecture is selected and expectations are not adjusted according to choices made. An example may be higher expectations on a virtual environment than is possible with deployed architecture.

* Not following best practices outlined in the APM Sizing and Performance Guide. One example is running all of the APM CE services on an undersized or 32-bit collector which can lead to various performance issues.

* The Architecture was sized correctly at the start. But the architecture is not reviewed periodically to keep up with changes in the environment.

My wish is not to see any of the above scenarios in future Support cases.

Time spent on architectural planning now may reduce or eliminate outage time in the future. Please consider the above for present and future deployments.

2.2 Another look at APM System Administration Philosophies (2014)

URL: <https://communities.ca.com/community/ca-apm/blog/2014/10/07/apm-blog-another-look-at-apm-system-admin-philosophies>

Introduction

Since 2009, I have been writing on APM system administration philosophy. This includes a Tuesday Tip, a recent Knowledgebase Article, and a section in an unpublished book. Why do I keep returning to this? Because I do not like seeing APM Clusters looking like below:



Frequently, support cases are opened where APM clusters are sluggish, inoperable, underpowered, and not providing full advantage of the solution. This is often the result of doing too little maintenance until it is too late.

System Administration Levels

The level of APM system administration is a continuum from minimal/as-needed operations to being proactive. This “choice” is often made simply due to workload rather than active deliberation. But this decision is an important one. The amount of attention provided to an APM Cluster impacts system stability, monitoring quality, frequency of upgrades, number of support cases, degree of false positive metrics, etc.

Once the System Administration approach is decided, then the associated efforts can be scheduled, planned, and staffed. One can evaluate periodically proactively or after outages if additional APM system resources and practices will help alleviate outages, improve system optimization and maintenance, and other benefits. As one progresses to the next stage of the System Administration continuum, additional tasks may be initiated. This may include creation of operational run books, monitoring strategies, associated project management, and more.

There is a common saying -- “cheap is expensive.” The same thing can be said of system administration inattention and benign neglect. It is my wish that this blog convinces at least one APM administrator to be more proactive as to the caretaking of their APM cluster by checking logs more frequently, upgrading to more current releases, and immediately eradicating found issues. A well-maintained house and APM cluster brings typically peace of mind. Here is hoping that the same happens for you.

References

Photo is courtesy of Karen Apricot and Wikimedia Commons.

Knowledgebase Article -

<https://support.ca.com/irj/portal/anonymous/kbtech?searchID=TEC617154&docid=617154&bypass=yes&fromscreen=kbresults>

Tuesday Tip - <https://communities.ca.com/message/98086748#98086748>

2.3 Four Keys to Successful APM Searches: Part One (2014)

URL: <https://communities.ca.com/community/ca-apm/blog/2014/07/27/four-keys-to-successful-apm-searches-part-one>

Introduction

"Nothing is hidden," while the Soto Zen founder, Eihei Dogen, was talking about "practicing life as a meditation", he easily could have meant searches for Application Performance Management information.

I have spent eight years searching continuously for answers to user questions for achieving stable and scalable APM monitoring infrastructures. And while others have been unable to find them after an "extensive" search, I have been to retrieve answers using the techniques listed below. It is my wish that this article will start an informative and wide-ranging dialogue on this topic. So let's talk about two different Search Strategies, shall we?

Clueless Searches

Tommy Sherman is under the gun to find answers to his questions. So he queries frantically just to get anything to appear in the results area. Because any result is far better than a blank screen. Many just flail away doing futile information searches with not much thought is given to the query. But this approach does not have to be followed.

Organized Searches

Garfield Lemay plans his searches before typing a single key. This includes the following:

Key #1: What type of query is it

Are you searching for:

- Answer to a simple question? (Information Request)?
- Steps to install/upgrade/configure an APM component?
- Causes of an Error Message or a performance/condition?
- Answers to multiple related questions? (Complex search)
- Release-specific/Site-specific information?

Knowing the type of query will determine

Which sources to search

- How long to search
- Results likely to be returned

Key #2: What is my APM Search Strategy?

The following are some successful search strategies to use?

- Start specific and get more general or Start general and get more specific?

The first approach goes for the “quick hit” and then expands the search. The second approach looks for analogous behavior in other components and then drills down. (An example of this is having a LINUX-based error. Seeing if this error occurs in other APM components and third-party products may result in more options.)

- Timeframe of searches. Some errors may be release-specific or occurred in the last two years. So more current information may be relevant. But some errors may have started with a few occurrences and later became more frequent. One example is the XML Parse Error in the TIM log. This used to happen infrequently. But with the arrival of SOAP, Web Services, and XML transactions, this became more common. So a wider search may show you the trends used to resolve an issue. And that gives you an idea of viable options and what was done before.
- Search terms, search terms, search terms.
A major reason for not finding information is simply the search terms used. Some guidelines for successful searches:
 - If there is a Java exception, searching on the offending class/method may yield a treasure trove of information.
 - Include the who, when, and where in search terms:
 - Describe the general operation (Such as install/upgrade, report, configuration)
 - Include the component name/version, third-party software version, and APM version. Again try this with one of these factors and then expand to all of them.
 - List what is and is not happening. (Such as no SLA reports)

Next time, I will continue on this topic. Until then, happy searching!

2.4 Four Keys to Successful APM Searches: Part Two (2014)

URL: <https://communities.ca.com/community/ca-apm/blog/2014/08/05/four-keys-to-successful-apm-searches-part-two>

Introduction

This second part of this series continues discussing the remaining keys for finding Application Performance Management information. I have found these four keys very helpful in getting answers quickly to various questions.

Key #3: Search in Installation/Configuration Files

One often overlooked area is reviewing the installation/configuration files for information on current and default settings. Some helpful directories and files include:

<EM-HOME>/config/ -- There are many more useful files here than can be covered. These include:

- tess-default.properties & tess-customer.properties. These files contain documented information about TESS settings for various APM CE operations.

- IntroscopeEnterpriseManager.properties. This includes basic and advanced Introscope Enterprise Manager settings and related annotations.

- tess-db-cfg.xml. - This file stores the APM database connection information.

<EM-HOME>\install\database-scripts

These are the many files used to install and upgrade the database. For example, create-tables-* and create-apm-tables-* (appmap fields) gives an overview of the fields and tables used in the APM database. This can be helpful in learning how to retrieve information for custom reports.

<EM-Home>\com.wily.apm.tess\com\timestock\tess\data\objects

I find the xml files invaluable because they show the fields used in various APM CE reports. An example of a daily statistical report xml file is StatsTranSetUserDaily.hbm.xml .

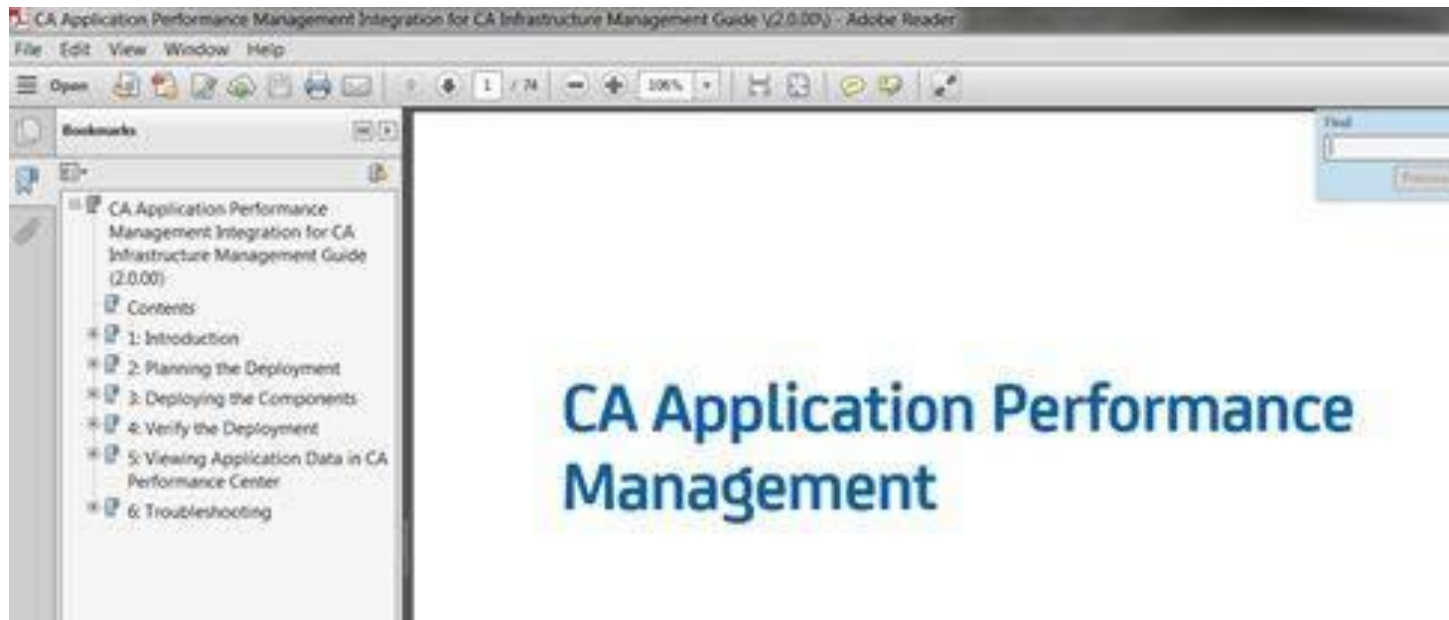
<EM-Home>\com.wily.apm.tess\WebContent\WEB-INF\jsp

This location is also helpful because it contains the web content used in creating APM CE (CEM Reports). And there are many more useful files/directories as well.

Key #4: Utilizing the APM documentation bookshelf

There are some customers that do not know about the full reader search capability in the APM documentation bookshelf. Go into any APM PDF file in the bookshelf, press Ctrl F, and then press the arrow in the blue box. You can then access the Full Reader Search and query search terms across the entire bookshelf. This can be helpful when a topic is covered in multiple

files/documents and when you also want to see the bigger picture on how common a search term is.



And if these four keys fail you, what should you do?

- Try a third-party web search site. These are very helpful for generic error messages for an application, operating system, database, etc. Many of the APM Community conversations are also searchable here.
- Ask around the Community message boards for suggestions and answers.
- Ask your CA Technologies Account Rep to ask internally for a response or a next step.
- Try to reproduce and experiment with a sandbox environment. Sometimes, doing what-if scenarios can reveal a lot of information. Making mistakes can be very helpful. My family learned just as much about playing chess from a computer program that made wrong moves as from one that was a stiff competitor!
- And if all else fails, open a support case. However, with more information being added each day to the APM knowledgebase, there should be less of a need to do that for many common basic and intermediate problems.

I hoped you enjoyed this series. As time allows, I will create other blogs on a variety of topics.

2.5: The Loneliness of the Solitary APM Administrator (2015)

URL: <https://communities.ca.com/community/ca-apm/blog/2015/02/08/the-loneliness-of-the-solitary-apm-administrator>

Introduction

There is a special place in my heart for the lone site APM administrator. Lack of time is always their challenge. They have often more than one product to support. The requirements of APM administration is always changing as more applications are frequently added to monitor. People are always coming up to them asking, "What does APM say about my application performance?" As a result, the environment may be subjected to bouts on instability keeping up with the increasing load. What can they do to get ahead? This article gives some possible suggestions

Mental Aspects

It is important to have a good mindset and not to feel discouraged when neck up with reactive issues. This includes the following:

- **Do your best today, do better tomorrow.** Keep at it!
- **Not knowing but keep going:** Attempt to move forward without knowing the final result. Each day things are changing for the better. Sometimes it take a long time for the seeds that you are planting to harvest. Believe in your actions.
- **Be Kind to Yourself in All Situations:** Sometimes in the middle of it all, we get "stress-stupid" and miss obvious solutions. Don't kick yourself for the rest of the day. Learn from the situation and move on.
- **Know you are never alone.** There are others that have been or will be in the same boat as you. Seek them out on the APM Community Site.
- **Enjoy the dance of APM Administration!** It is dynamic, challenging, and exhilarating.
-

Proactive Aspects

It can take some time to get out of reactive situation to a proactive state and an optimized environment. But doing the following, will get you closer:

- Take care of the "low hanging fruits."

- Clean your logs of minor errors.
- Make configuration changes that will have the biggest impact.
- Increase physical memory and heap size on MOMs and Collectors. Evaluate impact of this change looking at the perflog.
- Remove any unneeded metrics that could impact performance (such as socket metrics).
- Monitor only those applications and transactions that you really need to rather than those that are nice to have.
- Keep up with operating system, security, and third-party changes

- Undergo Planning Baby Steps

- Implement a monitoring strategy/implementation roadmap and review quarterly.

- Review architecture against current and future needs at least quarterly. Make changes before they are needed. Don't wait for "just that one more application" that will break things.
- Do frequent health checks, optimizations, and maintenance. This is not "do it once and forget it." (Just like dental cleanings.) **It must be an ongoing process.** Read the previous sentences again until that sinks in.
-

Knowledge Aspects

- With new releases every couple of months, the world of APM is fact changing. When evaluating a new release, one has to know which of their current problems are fixed , which features they are using are removed, and which new features to consider deploying. This is an ongoing process. Reading the known issues, release notes are an important way to do this
- Read APM Flip Board, Community, Tech Tips on best practices and field enhancements that may apply to you. There is a wealth of information therein.
- Take at least twenty minutes each day to build your APM knowledge.

Conclusion

If you keep up doing the above steps, you will see real results. Rather than pondering at the end of each day, "What did I do today", you will be excited with the small victories that you are seeing taking root. I would love to hear back from you on this topic. You can reach me at hallett.german@ca.com

2.6 The Wonder of APM CE Services (2015)

URL: <https://communities.ca.com/community/ca-apm/blog/2015/03/07/apm-blog-the-wonder-of-apm-ce-services>

Gone are the days of milkmen when you left numerous bottles outside your door. Almost magically in the early hours of the dawn back then, the old milk bottles disappeared only to be replaced with new ones.



The same sort of magic also happens with APM CE (CEM) Services. Each morning when arriving at work, there are new incidents and reports, updated statistical tables, and a cleaner database. (Note: If monitoring too many transactions, creating too many defects, or not following best practices on server requirements/architecture, then the above results may or may not happen.)

So let's talk about what is going on behind the scenes to get the daily report:

The **TIM or MTP/TIM** generates the needed ingredients of the recipe to make reports. This includes defects, statistics, and session/login mapping information (in cache). Looking at the timfiles log shows which files are or are not being processed. Most of the interim files are stored under /etc/wily/cem/tim/data/out. The stats files are kept for the previous hour while the other files are collected every five seconds.

The **TIM Collector** pulls the files off of the TIM. It also does the hourly or interval aggregation of statistical information. (This collector additionally does RTTM and Transaction Discovery processing. Because it is busy around the clock, a dedicated TIM Collector server is highly recommended.) If for some reason the TIM Collector gets bogged down in processing, then data may back up (But not lost.) on the TIM. A restart of the TIM Collector typically resolves this issue.

The **Stats Aggregation Collector Service** runs by default shortly after midnight. It aggregates the hourly into weekly, weekly into monthly, and monthly into yearly data. These processes run independently of each other. Up to two years of data can be processed. Depending how many defects that you have, your configuration settings, and whatever else that server is doing, this may take minutes or hours. Viewing the log in INFO not DEBUG mode allows viewing the processing over several days. Many of the settings for this are configurable in EM-

HOME\config\tess-default.properties\tess-customer.properties. (Or equivalent location. There are other files in that directory that can also be configured to impact performance.)

The **APM Database Cleanup Collector Service** is the least intensive service and typically runs on the Stats Aggregation Collector. At three A.M. it kicks off vacuuming and soft deleting old records and tables.

The **APM Database** stores the output from the TIM in Postgres and Oracle database tables. If available, it also stores Appmap data. Without proper maintenance, it can grow quite large and can impact performance in displaying reports. Typically large tables are users, user groups, session maps, app maps, and defect metadata.

The **MOM GUI** is used to display and to schedule the creation of reports. But it is dependent on all the above servers to do their jobs before a result is shown. So keep in mind all of the magic involved when looking at your next report.

Thanks to Safekidda and Wikimedia.org for use of the image.

2.7 Problem Solving: Six Ways to Avoid the Wrong Path (2015)

Introduction

In Colonial Williamsburg near where I live, there is a maze in the garden behind the Governor's Palace. At the start of the maze, the sign invites you to go left. Once inside, you will get easily lost and misdirected. But eventually you will find your way out.

At the end you discover, if you went right in the first place, you would already be by the exit.



Sometimes, I see similar behavior in APM problem solving where if only someone asked just one more question or did just one thing differently, then the issue would be resolved already.

Six Ways not to get Trapped.

Here are six ways to avoid being misled from least to most importance.

1. Ignoring False Positives

Sometimes log messages look like warnings or errors but are really information messages. This is especially true when turning on DEBUG or VERBOSE mode in Introscope logs.

Knowing if something is really an error worth pursuing or not can help maintain proper focus.

2. Removing Incompatible software

Often, an issue is found due to software incompatibilities. Such as an EM release being later than an agent release. Or a JRE release that is unsupported in the current APM release but is supported in a later one. Checking the compatibility guide is a quick task to do. Please do this early not later with a problem as to not to delay resolution.

3. Find out when the issue started.

Often, customers have a good idea when a behavior or error began to manifest. However, what they might not be aware of is the issue started long before that. So reviewing logs may get to the real state date. Once there, the behavior may be tied easily to some event.

4. Discover if something is a symptom or a problem.

Many issues may be a symptom of something else such as performance issues, reaching server capacity, unoptimized software and more. Peeling the layers of the onion to get to the real issue may take time. But in the long run will bring about system stability and scalability.

5. Not checking if this is a Known Issue

A five minute check of the release notes for your current and future releases, as well as the always updated knowledgebase may quickly indicate if something is a known issue or not. This will help avoid hours of aggravation.

6. Determine if there are Cascading Issues

Recently, there was a customer had an issue where data was not leaving the TIM. So, the focus was placed there. However, the real culprit was the Tim Collector. It was not doing its job since it was performing other services (A common issue). As noted in the Tech Tip below, knowing the functional workflow of a task including likely failure points and what is next in a successful workflow can be very helpful.

<https://communities.ca.com/message/241725272#241725272>

I hope this is helpful and aids you in avoiding future issues. Please respond to this blog as to other things that can misdirect a problem resolution effort.

2.8 APM Administrators: A Dying Breed? (2015)

Introduction

It is 2045 and the following announcement is made by some software company:

And with the release of APM Version 235, the software is now 100% self-aware. There is no need to have administrators anymore. Our software can and will do it all for you. Contact our sales staff to see a demo today.

Is the above strictly a figment of some veteran science fiction writer's imagination? Perhaps not. Google and others have started baby steps along this path. While some worry about [robots taking their jobs or the rise of nano-robotics](#), I think about what a world with truly-self-aware software would be and is there still a role for APM Administrator.

Self-Aware Software

Earlier this year, I noted [all the things that a good APM Administrator would do](#). I would expect truly self-aware software do the following and a lot more.

1) Install and Upgrade Flawlessly

Installations/Upgrade/Migrations would have no problems from start to finish. If needed, virtual servers are created and configured. The correct combination on configuration parameters are determined to handle current and the anticipated needs for the next six months or so.

If an issue is found, then it must be self-corrected immediately. After the installation, the software monitors for exceptions and changes in metrics that deviate from the pre-installation baseline.

2) Self-correct/Self-configure

The software tirelessly and continuously monitors logs, supportability metrics, etc. for limits that are about to be approached and changes the parameters to a new optimized setting. (Which would also keep in mind the current settings of other configuration parameters.) If needed, conversations are initiated with the self-aware network and application server software to resolve issues.

3) Conduct Long-term Analysis

The software also looks periodically at historical trends and uses this information to make better decisions. This may include determining time periods where the monitoring software is stressed more than usual.

Wrap-up

So what do think? What other characteristics of self-awareness would APM software have?

In conclusion, I don't expect this very soon. After all, we can't even do a good [grammar checker](#) or [translator](#).

Note:

This above blog should not be seen as an endorsement of any software. No future plans about any APM release for CA Technologies or other software companies are contained herein.

[2.9 The Confessions of a Serial Change Agent: Part 1 Overview](#)

[This is a two part series on being a change agent. This first part is the characteristics of being a good change agent.]

Introduction

It is Sunday morning somewhere. And someone is frazzled after a long sleepless night where they are half afraid of making the wrong choice and the other half consumed by the implications of making the correct one.

In these situations, there are three real choices:

1. Accept the situation. Nothing changes and you have to grin and bear it. One becomes stuck in time where they are not moving things forwards or backwards. Health implications may arise. Robert E. Quinn talks about the "Slow Death Dilemma" because of the reluctance to deal with things.

Like American journalist Esther Wojcicki says, "The worst thing that you can do in life is whine about what you can't change."

2. Leave the situation. By doing this risks that things will be about the same or better. But there is the other scary possibility the situation becomes far worse

3. Shake up things by being a change agent.

Ask yourself this question:

If you are not going to try to improve the world now then when?

The change agent sees no obstacles only unexpected gifts.

Why be a change agent?

I have been a serial change agent for most of my life. In each point in my career, I have left successfully the situation better than I found it.

This is what I said about being a change agent in my book *Musings*:

Being an "in the trenches" change agent is a huge risk. You are the pioneer who gets the arrows in the back. You have no official responsibility to do anything. People you don't even know may yell and scream at you because they are frightened by your efforts. And if you succeed, others likely will take the credit and build on what you did.

But what about the rewards you may achieve?

These include the following:

- *Making a real difference in the world by overcoming the conventional wisdom and changing something for the better.*
- *Personally bringing to the world happiness, increased productivity, and moving things forward.*
- *It is the greatest self-reward in the world for having done this.*
- *It gives you energy for undertaking the next major change.*
- *You will learn things about yourself and your capabilities that you did not know already.*

Characteristics of a Change Agent

These are the behavior patterns of a successful change agent:

- Fearlessness
- Active Listener
- Good Communicator
- Can clearly explain the difference between the AS-IS (Current) and TO-BE (Future) environments.
- Empowers themselves to make needed changes.
- Works with the end in mind.
- Has a strong sense of commitment and incredible perseverance.
- Understands well their own strengths and weaknesses.
- A good facilitator
- A talented builder of ad hoc and permanent coalitions.

Characteristics of a Serial Change Agent.

The most involved change agents move from one situation to another doing what they can to improve the situation. Let's call them **serial change agents**. In addition to the

above, they do the following:

- Have a strong sense of perseverance
- Ability to recharge on the fly
- Believes in the reverse of the Butterfly Effect -- if you want to change one thing, change everything.
- Uses some of the techniques learned from previous change efforts.
- Dynamically learns and adjusts.

I would love to see conversations spring forth on being a change agent in and outside CA.

I want to hear more what you have accomplished after seeing a need for improvement.

You can reach me to discuss further at hallett.german@ca.com

Next time, I will talk about some of the things that I have done as an APM change agent.

References

<http://www.inc.com/magazine/201510/jeff-bercovici/the-dna-whisperer.html> Where the no obstacles and the Esther Wojcicki quote came from

Quinn, Robert E. *Deep Change: Discovering the Leader Within*. New York: Jossey Bass 1996 Print. My favorite book on the subject.

<https://www.smashwords.com/books/view/399500> Musings. Includes my very incomplete look at this area.

[2.10 The Confessions of a Serial Change Agent: Part 2 APM](#)

Part 1 can be found at [The Confessions of a Serial Change Agent: Part 1 Overview](#)

Introduction

This is the second half of a two-part blog discussing the rarest of creatures -- the serial change agent. Last time, I covered the characteristics of being an excellent change agent.

This month, I will talk about some of the things that I have done as an APM change agent. I will start from an individual level and work up.

Seven Steps

1. Ask questions

Never stop asking why things are the way they are. Try to understand the various challenges and skill sets of key stakeholder organizations and players. This helps to determine the best way to make changes. The [KB on MTP and TIM compatible releases](#) was one such recent result.

2. Identify gaps and do something about them.

A wise manager at CA said to me, "Anyone can point out something's wrong in a situation. Far fewer hear the call and try to resolve things." An example of the former is the fictional character from Saturday Night Live: [Oscar "Fixit" Rogers](#). I have taken my ex-manager's words to heart. If I see something is not correct, then the next immediate step is an actionable item for improvement.

3. Document like crazy.

Since 2006, I have documented through emails, [Tech Tips](#), [Tech Docs](#), knowledge documents, blogs, internal presentations, and cases the various best practices in APM to produce stable, scalable environments. This sizable body of knowledge has enhanced the APM administration experience.

4. Look for like-minded individuals.

They are the ones that raise key issues and then do something individually about it. Keep talking with each other and good things will happen.

5. Form ad hoc project teams to get things done.

With #2 in place, small teams can be created to quickly implement needed change deliverables. Then when the work is done, end the team. Repeat again when the next opportunity forms itself.

6. Form "permanent" teams as needed.

I took over running the APM Documentation Board which meets each month. This group has quietly improved long-standing documentation shortcomings and works hard to make APM technical documentation better suit your needs. Examples include adding the [APM Troubleshooting](#) and [Error](#) sections, providing a link to the [MTP documentation on span ports](#), and much more. [Have an APM documentation concern and want to make sure the right people hear about it?](#) Feel free to drop me a note at hallett.german@ca.com

There are other efforts as well that may be a subject of future blogs.

7. Address concerns across products.

From my experience with the Problem-Solving , Knowledge Management, and Search Engine cross-product projects this year, I have seen how conversations and linkages across products leads to better customer experiences through streamlined processes.

Conclusion

In summary, change what you can today, change even more tomorrow. **The rapid pace of transformation in the world means that we can no longer afford to be comfortably on the sidelines.** You are going to be here for some time. So, might as well make a difference. no matter how small. I hope these two articles inspire some one to get out of their comfort zone and help transform their universe in some way.

Section 3: Tech Tips (2015)

3.1 APM CE Knowledge Levels

URL: <https://communities.ca.com/message/241730878#241730878>

Introduction

Over the years, I have written about APM CE administration best practices including administration philosophies and recommended tasks. But there has been one area that I have ignored purposely. Part of the administrator's day is assessing current knowledge levels and determine the needed training. But how does one do this self-assessment?

Knowledge Levels

Some years back, I created a list of what I thought were APM CE knowledge levels. This was released internally but not publically.

The knowledgebase article below provides this list:

<http://www.ca.com/us/support/ca-support-online/product-content/knowledgebase-articles/TEC1590953.aspx>

What are these various levels are based on? It includes the following:

- Knowledge to adequately perform basic APM administration daily and weekly tasks.
- Understanding of and ability to perform integration tasks at various levels
- Experience in planning/performing trouble-free APM installs and upgrades.
- Familiarity with using/producing all types of canned and customization reporting.
- Strong comfort level in customizing APM through such things as evidence collection, HTTP Plugins, Custom PBDs and Dashboards, and other means.
- Efficient and effective troubleshooting for one or more APM products or CA Technologies integrations. This may require some idea of how these products work behind the scenes, knowing how to tell a real error from a false positive one, understanding the context behind a problem, and using tools and techniques to optimize and troubleshoot a situation. This requires a wide range of skills such as reviewing LDAP authentication sessions, understanding database SQL queries in logs, and diving into network packet captures to see if traffic is being filtered out.
- Ability to do architecture planning, cluster optimization and maintenance

How to improve your knowledge level.

So how does one improve their APM knowledge levels? Through the practice of consistent effort. Here are some suggestions below:

- There are variety of documents, trainings, and electronic forums available to improve your breadth and depth of APM CE knowledge. Become familiar with them.
 - Build yourself a training plan and try to get something completed each month. Reevaluate the success of your training plan and periodically review what should be included/excluded.
 - Ask questions in forums and training sessions. One good answer can improve our understanding and help us succeed.
 - Take the time to try out various scenarios in new and current releases in a laboratory environment to see what is possible.
- Sadly, some sites still do not have test environments for this purpose.

Questions for Discussion:

Do these assessment levels make sense to you? What do you like and what would you change?

1. Do you have a training plan that you follow and occasionally modify? What have you learned from this process?
2. Do you have one unanswered APM question that you wish to have a clear and definitive answer for? If yes, what is your question?

3.2 Planning to Monitor your Application by APM CE

URL: <https://communities.ca.com/message/241734832#241734832>

Introduction

It is the fifth straight year that I have been writing Tech Tips. I have learned so much on the journey from interacting with many of you and want to convey my thanks.

This is a special tech tip since it is number fifty. [It could be number forty-nine or fifty-one. But close enough :-)]

In previous tips, the Transaction Definition Lifecycle was covered in detail. But one area was left out and that is the planning phase. This is covered in part in a new KB:

<http://www.ca.com/us/support/ca-support-online/product-content/knowledgebase-articles/TEC1290451.aspx>

But this KB does not include an overall framework. A modest start of doing this is provided below.

First Step: What and Why am I Monitoring?

One or more applications are to be monitored as part of a company's application monitoring strategy. Applications may be prioritized based on relative health, political concerns, application complexity, and business/financial focus. It is important to articulate why you are monitoring and make sure your implementation plan follows that.

This can be due to reasons similar to the above selection criteria:

- Desire for high customer satisfaction.
- Lost revenue concerns.
- Improve/increase application availability
- Eliminate potential problems for new releases

The above reasons will determine:

- The business transactions/transactions to be monitored.
- The metrics to be monitored
- The reports to be generated.

Second Step: Planning the Business Transaction Recording Lifecycle

Once the what and why is completed, then the traditional project management planning stage is initiated. This includes

- What are the roles and tasks for the various lifecycle processes?
- Who will monitor the monitoring for overall success?

- Will reporting used reactively and/or for capacity planning? For defect forecasting?
- How integrated with APM CE Monitoring be with the application development lifecycle?

Take the time to do this step right including capturing assumptions, risks, using dedicated resources, have clear tasks, and realistic timeframes.

Step Three: Know your Application and Diving In.

The better you understand the application, the less likely that you have issues. This includes the following:

Complexity of application	
Simple	HTTP/HTTPS Request-based traffic, ASP/ASPX
Intermediate /Complex	XML/Web Services, FLEX, JSON, SAP, Seibel, Microsoft Web Applications, Oracle Forms, etc.

- Infrastructure setup (Web servers, load balancers, firewalls, reverse proxies, network switches)
- Location of TIM in relation to the application's server.
- Application's use of SSL (cipher suite TLS version), and other protocols.
- Compliance to HTTP Protocol (includes content-type, accept language).
- Verify round trip affinity (The transaction request and response goes through the same set of servers.)
- Sitting through walk-throughs with the developers on key transactions including determining how to identify users/sessions, group users, key parameter-value pairs to identify transactions, possible overlapping transactions.
- Go through several dry runs on recordings/monitoring to see if what development said above will work or if another HTTP field needs to be added to the application.

Thank you all for your support and feedback. The quest for promoting/achieving scalable and stable APM infrastructures goes on. I plan to keep monthly tips for as long as I can.

Questions for Discussion:

1. Are you doing the above planning process? If so, please share the details.
2. Which one of the fifty tech tips was your favorite?
3. What other topics would you like me to cover?

3.3 APM CE Technical Tip Overview

URL: <https://communities.ca.com/message/241759173#241759173>

<https://communities.ca.com/servlet/JiveServlet/downloadBody/231152516-102-1-14192/Overview%20Tech%20Tip.jpg> (diagram)

Last week I received this note.

Dear Hal:

I love all of your Tech Tips. However I want to know the big picture that ties this all together. Or does that image not exist?

Dying to Know,

Tom from Tewksbury

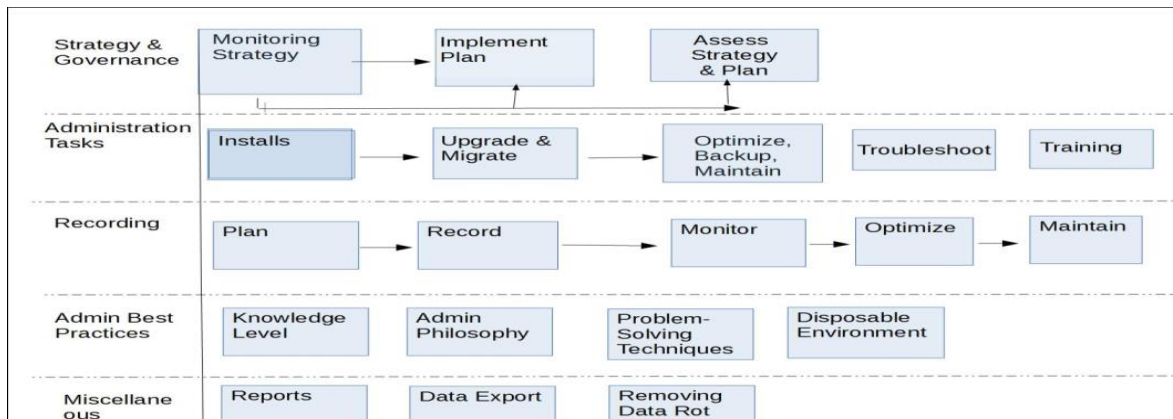
Dear Tom:

Thank you for your note. I hope you are having a mild winter. Fortunately for you, there is such a big picture and I am sharing this month with you and other users. Let me know if meets your needs.

Thanks,

Hal German

As I just told Tom, all of my 50 Tech Tips do tie together and the attached diagram at shows how. If you read/re-read these articles, please keep this diagram in mind.



Questions for Discussion:

1. Does this diagram help give a clearer view of the APM CE big picture?
2. Are there other topics that should be included?
3. What other topics should I cover in the future?

3.4 Tim Settings Part 1

URL: <https://communities.ca.com/message/241784124#241784124>

Introduction

One of the overlooked areas for TIM problems are TIM settings. This underdocumented area will be explored below and in the articles in the next few months.

Overview

The TIM has a variety of Global settings that have the format of names and a value. The name is usually two parts such as BtStats/Enabled where the first part is a prefix that describes a class, function, or category. The second part is the operation or subcategory to be configured. Note that there are one-part setting names such as MaxLogSizeInMB.

Some settings may not be explicitly defined so will use default implied settings.

Some of the common Name Prefixes include:

AutoGen	- Transaction Discovery
BtStats	- To enable/disable RTTM, change log level and other functionality.
DiskSpace	- Disk space settings.
Flex	- Flex Processing
FlexAMFParser	- Flex Processing
Parallel	- To enable using workers for MTP.
SetTimTime	- Set the Tim Time maximum difference.
SiteMinder	- SiteMinder settings such as trace enabled.
TessTimeZone	- Hour Offset
TraceFilters	- TIM Trace options
TransactionInspection	- TIM Transaction Settings options
XmlPlugin	- Enabling XML Plugin

These settings are accessed or changed through the *TIM System Setup* > *Configure TIM Setup* page via your browser. Once there, you may view the current settings and then change them. Or you may add a new name/value page. After the change is made, a TIM restart is NOT needed.

Besides looking at the GUI, there are other ways to view the TIM settings:

- 1) When collecting the TIM logs and configuration files from the TIM system setup there is a settings file in the status folder.
- 2) Using configtool (which is a python script) allows you to read, write, or delete one or more entries in the tim database. It is found in the `/etc/wily/cem/tim/bin` or `/opt/CA/APM/tim/` (APM 9.6 and later) directories. There are various command line arguments to do this:

- f -- the Tim settings filename and location. This is required. Filename is timsettings.db
- g -- display specific name/value pair.
- l -- list all values

- n -- create a new TIM settings database
- u -- unset a value
- F -- Fast - Does not tell TIM of change. (Normally TIM is aware of all changes made to the TIM Settings database.)

Examples:

1. Getting all settings -- /etc/wily/cem/tim/bin/configtool -f /etc/wily/cem/tim/config/timsettings.db -l
2. Getting a specific setting -- /etc/wily/cem/tim/bin/configtool -f /etc/wily/cem/tim/config/timsettings.db -l | grep 'Tess.*IpAddr'

3) On TIM Restart, all of the TIM Settings are listed included those not shown in the GUI. The difference is this also includes the implied TIM Settings. So while a TIM Restart after a change is not needed, it is recommended to see if the TIM setting change took place. The line will look something like this: (This is for a MTP so there is a worker number.)

```

Mon Mar 30 09:24:32 2015 18717 w2: ComponentMonitor: DB: BTStats/Level = 0
  ^           ^           ^           ^
  Date &Time           TIM Process           Setting Name Setting Value

```

The data is stored in a Berkeley database which has nothing other than the TIM Settings name/value pairs. The TIM process does not need to be running to access the database.

Notes:

1. Some programs will write the TIM Settings to the APM database for you. (nCipher is one such example.)
2. There is no database name/value validation. So please check that your name/value entries are correct after entry.

Next time we will talk about which settings are documented and where. the final article will be on troubleshooting

Questions:

1. Which TIM Settings have you changed that solved problems or improved performance.
2. Is there a particular TIM Setting that you want to be better document?

3. Please mail to hallett.german@ca.com future topics to cover

3.5 Tim Settings Part 2

URL: <https://communities.ca.com/message/241790546#241790546>

Tech Tip: TIM Settings Part Two

Introduction

This is the second of three parts on TIM Settings. TIM Settings traditionally have a one- or two-part name and value. If not listed, then the default values are used. A list of TIM Settings can be seen in the TIM log on startup.

References:

At this point, there is not one place that you can see the TIM Settings. Below is an incomplete list of TIM Settings from the APM documentation . This is from the 9.7 documentation.

- 1) The new *APM Troubleshooting Guide*. This includes information about the
 - Flex and XML Plugins which are enabled/disabled through TIM Settings.
 - IgnoreOnlyRequestErrorComponents setting to help limit the number of missing response defects.
- 2) *Transaction Definition Guide*, there are the directions to disable the XML Plugin
- 3) *Installation and Upgrade Guide*, covers changing TessCollectorIpAddr or the TessIpAddr values as needed.
- 4) *Configuration and Administration Guide*, lists the following TIM Settings :
 - RTTM or also called BtStats. Notethat enabling/disabling this functionality in APM 9.7 has been moved to a new file called btstatsconfig.xml.
 - DiskSpace/MinFreeDataSpaceInPercent for freeing up TIM disk space
 - ProxyForwardHeader to set the real Client IP address for defects.
 - TessCollectorIpAddr or the TessIpAddr to set the IP addresses for MOM-TESS and Tim Collector. It also talks about MOM failover and TIM with these fields.
- 5) *Security Guide* has directions to do the following:
 - Set the MaxDefectRequestBodySize to increase/decrease the request body size.
 - Set timTessCommunication.useSsl to enable/disable TIM-TESS EM SSL communication. This at one point was an undocumented property.
- 6) *Sizing and Performance Guide* includes the following settings:
 - DiskSpace/MinFreeDataSpaceInPercent which was already discussed

In the CA APM for Infrastructure Management Guide talks about

- Parallel/UseWorkers. This is an important setting and only should be enabled for MTP and

disabled for non-MTP TIMs with or without Napatech cards. If misconfigured, this can dramatically impact TIM performance.

Implications/Other References

This also means that there are whole areas of TIM Settings not in APM documentation

- SiteMinder settings (I will bring back the KB on this soon.)
- Trace Filters
- Transaction Inspection
- Miscellaneous settings (MaxLogSizeInMB, SSL Ageout, others)

There are 13 KBs on TIM settings. These include:

<http://www.ca.com/us/support/ca-support-online/product-content/knowledgebase-articles/tec606479.aspx> -- MaxLogSizeInMB

<http://www.ca.com/us/support/ca-support-online/product-content/knowledgebase-articles/tec603778.aspx> -- MaxDefectRequestBodySize

<http://www.ca.com/us/support/ca-support-online/product-content/knowledgebase-articles/tec604032.aspx> -- 403 Error and TessCollectorIPAddr/TessIPAddr

Notes:

1. Not all TIM Settings are documented due to a variety of reasons.
2. Documentation may include new but not deprecated or removed TIM Settings.
3. There is generally little discussion of what are the performance implications of raising a value past the default.

Questions for Discussion:

1. What would you like to see me cover with TIM Settings?
2. What are areas of TIM settings that need to be documented?
3. What should be covered in other topics?

3.6 Tim Settings Part 3

<https://communities.ca.com/message/241800601#241800601>

Introduction

This is the last article on the underdocumented area. The previous two articles gave an overview of and provided references for TIM Settings.

This last section goes into Tim Settings troubleshooting.

The previous sections can be found below

<https://communities.ca.com/message/241784124#241784124> - Part 1: Overview

<https://communities.ca.com/message/241790546#241790546> Part 2: References

Problems with TIM Settings typically fall into these areas

- 1) Invalid Tim Setting name or value.
- 2) Memory issues/TIM restarts caused by TIM Settings.
- 3) Communication issues between TIM and Tim Collector/MOM.
- 4) SSL-related issues due to TIM Settings.
- 5) RTTM not working.

Let's look at each of these in turn.

1. Invalid TIM Settings or value

Creating or updating a TIM Setting is a manual operation. In some cases, the TIM Setting has a default setting even if not explicitly set. So to be changed, it also needs to be manually added. Because of this, it is far too easy to configure an invalid Tim setting name or value. As far as I can determine, not all names/values are extensively validated at this time so be very careful in doing this. If there is any doubt, then use the configtool previously mentioned to see what is stored in the database.

2. Memory issues/TIM restarts caused by TIM Settings

Under load, TIMs can use more CPU or memory when processing data with HTTP Plugins. Some of these built-in plugins are turned on by default but not explicitly set. Only by using the configtool or viewing the TIM Setting after a TIM restart can you tell that this is the case. The two most troublesome are

Flex

When you are not using Flex, then disable them by setting the following properties to 0:
FlexPlugin/Enabled
FlexAMFParser/Enabled
FlexAMFXParser/Enabled

XML

XmlPlugin/Enabled

Another reason for this is enabling Parallel/UseWorkers when a TIM is not a MTP or disabling this setting when a TIM is a MTP.

3. Communication issues between TIM and Tim Collector/MOM.

This is a very common issue. The culprits are typically one of two settings

TessCollectorIpAddr = 10.0.0.41
TessIpAddr = 10.0.0.36

Typical problems are due to the following:

- An invalid IP address.
- A valid IP address but either the TIM or TIM Collector/MOM cannot reach it.
- Port is blocked.
- EM has two IP addresses and TIM is attempting to connect to an unreachable one.
- EM is not running

In any case, this is some sort of communication issue which involves changing the configuration value or opening a server port.

Until that is done, you may see various 400 or 500 errors on the Monitors tab in the APM MOM CEM GUI.

4. SSL-related issues due to TIM Settings.

SslSessionAgeOutCount = 100
SslSessionAgeOutSeconds = 300
DisableTLS11And12RecordsProcessing

Depending what these are set to, you may see many SSL Session Timeout and since APM 9.1.7 TIM SSL bad data (TLS 1.1/1.2) records messages in the TIM logs.
(The latter is true only if the TIM trace option for SSL is enabled.)

5. RTTM not working

BtStats/Enabled = 1

Enabling or changing other settings may increase TIM log size, impact recordings, and users being identified (proxyforwardheader), etc.

But these are minor compared to the other issues mentioned previously.

I hope this series has been of used to you and you further explore this area.

Questions for Discussion:

1. Was this a useful series?
2. What are areas of TIM settings that need to be documented?
3. What should be covered in other topics?

3.7 E-Commerce and User Groups

URL: <https://communities.ca.com/thread/241695211>

Introduction

We continue looking at APM CE setups that may be useful but can have hidden costs including opening support cases. This week we look at various user group situations.

Some APM administrators may not realize that they will get user information in defects whether it is set for enterprise or e-commerce mode. The difference is enterprise mode supports a wider range of user group identifiers and user information is not kept for historical reports. (Ecommerce only supports Request attribute user groups and Unidentified Users.)

Here are some common scenarios that can cause issues:

Scenario 1: Using case-sensitive user names/High number of Inactive Users

By selecting the case-sensitive user name setting, Gwashington,gwashington, and gwashingtonN are seen as three different users. This may be valid user names or user typos. For a busy enterprise mode application, this can fill up a database quickly. Combine this with a low setting for the number of days that a user becomes inactive and your APM database can fill up quickly. So cleaning up the database and changing the case-sensitive settings are two ways to alleviate this. Another way may be setting to e-commerce mode or removing the user identifier.

Scenario 2: Using a User Identifier with Ecommerce

The documentation states the following when having many user groups due to request attributes or many Client IP Subnets:

Important: If your installation is a large-scale e-commerce site (hundreds of thousands or millions of users) use caution with this feature. The number of user groups created can become unmanageable.

In spite of the warning, some customers utilize a user identifier in e-commerce mode. This results in the following:

- Many users being created and stored in the database.
- EM busy collecting user and session information which gets written to the database each hour. This eventually will fill up the database and it will become sluggish

So deploy users/user groups in ecommerce mode with caution and periodically assess the need for periodic database maintenance.

Scenario 3: Exceeding the Maximum APM CE User Group Limit

In the TESS-default.properties, there is a setting called maxUserGroups. It has a default value of 5000. Once reached, users are placed in the New Users Group and additional user groups are not

created. However, the documentation is silent on what the highest that this value can be configured. Exceeding 5000 user groups should be done with care and will depend on your architecture and EM/database hardware.

Users/User groups are a great APM feature. But its impact on APM responsiveness should be reviewed as an ongoing task.

Questions:

1. Are you a site that has seen one of these three scenarios? What did you do to resolve the situation?
2. If you are using e-commerce with user identifiers, is performance and functionality working as expected?
3. What other topics would you like me to cover?

Addendum from Lynn Williams:

NOTE: As of 9.1 the **maxUserGroups** clamp property in **tess-default.properties** has actually moved to the new EM config file **apm-events-thresholds.xml** with property name **introscope.enterprisemanager.max.transaction.user.groups** : [Manual Configuration File Upgrade Overview](#)

[3.8 APM Tech Tip: Going It Alone: Strategies for APM Implementations/Administrations When Lacking Key Resources.](#)

Introduction

This was inspired by a prior comment in the Community. Your networking, database, and application groups are key for a successful APM deployment or resolving APM usage issues. This is especially true for APM CE (CEM). However, sometimes these valuable resources are not nor ever will be available. This is the result of a variety of reasons. This article talks further about this situation and how to work around it.

Causes

An APM administrator can fall in this situation because of various outcomes. These include but are not limited to the following:

1. Politics or Company Culture

In some businesses, numerous silos exist and one group can only contact another in rare circumstances following rigid procedures. The lead time for doing this is typically very long.

2. Availability

These groups have their own priorities and projects. So they may not have the bandwidth to assist. When they do have a spare cycle, it may be too late and implementation may too far along. If they do not like how things are being done, this can delay a project even further to retrofit their requirements.

3. Fear

Whether it is reluctance or possible loss of face. This may have been due to having gone various times in the past believing there was a problem and the appropriate team showed this was not the case. (Also known as the cry wolf scenario.)

4. Security

Other groups have the access to login, need to perform a pre-implementation review before giving the approval to go ahead, or some other reason. However, it takes times to get them engaged.

Possible Resolutions

Overall

1. Learn to become self-sufficient as much as you can. Find existing documents on your intranet that may help you such as network diagrams, application design documents, database implementations. It never hurts to ask for copies if you are unable to find them.
2. Work on breaking the silos and building relationships. Show how the business depends on successful and optimized application and application infrastructure monitoring. Accept that this will take time and there may be setbacks along the way. But taking a long term view will ensure eventual success. Having an executive champion behind you helps facilitate things faster.
3. For APM CE, run a periodic transaction discovery recording to learn about new or changed applications/transactions. Some customers use transaction definition catch-all rules to do the same thing.
4. Ask for increased access rights especially during installations/upgrades.
5. For outages and lesser monitoring gaps, do an after-action review and invite these same groups. Build trust and buy-in through their involvement. Publishing the results shows how ALL parties can better improve their responsiveness and effectiveness by just reviewing the facts in a non-judgmental manner.

Networking

1. Start using powerful tools such as Wireshark to do your preliminary analysis before contacting a networking team.

Application

1. Get their buy-in on the monitoring strategy and APM deployment roadmap. Once they see the value, they may make themselves more available.

Database

1. Start using database tools to perform queries reviewing table health. You may have to request from the DBA to have read access. The results is you can better and faster assess APM database operations and health.

Questions for Discussion

1. Has one of those scenarios happened to you? If so, please feel free to share your experience
2. How did you resolve the situation?
3. What did you do to ensure that the same thing would not happen in the future?

3.9 APM Database Corruption

URL: <https://communities.ca.com/message/241823771#241823771>

Introduction

What does APM database corruption really mean and what can I do about it?

Database corruption is a concern for APM Administrators. It can range from the database being completely inaccessible to part of a record being partially destroyed. This is a brief introduction to this area and what to do about it

Overall

APM databases support various releases of Postgres and Oracle. Both of these can be subject to database corruption. Typical reasons for this include

- Hardware failures including disks, Raid controllers, CPU, etc.
- Operating System issues including bugs or corruption.
- Third-party software such as virus checkers, security hardening, etc.
- Database Misconfiguration or reaching database limits
- Other factors (Such as operator mistakes.)

Troubleshooting

Step 1: Determine Scope and Extent of Problem with APM Database

Before contacting Support, it is always helpful to better understand what happened. This includes:

- When the issue first happened
- The scope of the issue -- Is it the entire database, an index, or one or more tables impacted?
- Did the database run out of disk or table space?
- Are there error(s) and Java Exceptions in the database and APM logs that may help understand what the issue may be?
- Checking when my last database or Configuration/business transaction export is from.

Step 2: Backup the database directory, Do a Configuration/Business Transaction Export to save what you can.

- Grab a copy of the database and configuration as soon you can. This includes
 - * APM database backup via command line or pgbackup
 - * APM Configuration Export
 - * Screenshots of APM CE Configuration
 - * APM Business Transaction Export
 - * Screenshots of Appmap in Investigator
 - * Save individual APM database tables

Step 3: Open a Support Case

- Open a support case reviewing what was discovered in Step 1 and 2
- Upload a copy of the database so recovery scripts can be created by Engineering.
- Test the recovery scripts on a non-production database and assess results
- Then deploy recovery scripts on production database and assess results
- Use database recovery tools as directed by APM Support
- Run an APM vacuuming as needed.

Step 4: Proactive Steps

In the future consider doing the following:

- Evaluating all the applications, defects, metrics being monitored by APM to see if really needed.
- Raise defect thresholds
- APM vacuuming runs each night but do additional vacuuming as needed
- Periodically check APM database size and investigate reasons for increase
- Perform frequent APM backups and Exports

Additional Information:

Database Maintenance and Administration References

<https://communities.ca.com/thread/116071656>

Postgres

<http://thebuild.com/presentations/worst-day-fosdem-2014.pdf> Presentation on Postgres database corruption (Non-APM).

<https://wiki.postgresql.org/wiki/Corruption> Postgres & Corruption
https://bucardo.org/wiki/Check_postgres Nice free utility to check Postgres.

Oracle

<https://oracle-base.com/articles/misc/detect-and-correct-corruption> Tools to repair Oracle databases

Discussion Questions:

1. Are there other things that should be added as steps to do?
2. Are you using Database Failover/High Availability solutions? If so, please share your experience
2. What other topics should I cover?

[3.10 Ten Central Myths of APM Administrators](#)

URL: <https://communities.ca.com/message/241831677#241831677>

APM Administration: Ten Central Myths

Introduction:

For religious scholars, the term central myth refers to key propositions about that faith which is often reinforced through the retelling of a story supporting the myth. In recent days, this has come to mean beliefs that are untrue. In talking to many APM administrators, they seem to support various implied beliefs that are left unsaid. Adherence to these beliefs can lead to time-consuming tickets, and unexpected outages.

Below are ten common APM administration central myths and what they could be replaced with.

Ten Central Myths:

Central Myth #1: APM Administration is "one and done." All I have do is install/upgrade the product, assist in adding more applications to monitor, and just make sure things stay up.

New Myth: APM Administration is far from a "one and done" proposition. Proactive maintenance and optimization ensures a smooth running system.

Central Myth #2: APM cluster architectures do not need periodic reviews. Only look at increasing server capacity or number of servers only after things break.

New Myth: Periodic architecture reviews at least 2-4 times a year can handle increases in demand as well as factor in successfully future scaling needs.

Central Myth #3 I can go with the default APM settings. This should handle my needs for some time

New Myth; APM default settings work in many cases. But they often make certain assumptions that may not apply to your environment. Fast growing environments need frequent configuration parameter review

Central Myth #4 No need to clean my network traffic for TIM. It can make do with what I provide

New Myth: Unfiltered network traffic may lead to higher decode failures (because of TCP false connections) and make the TIM/MTP work harder in processing the extra traffic.

Central Myth #5: I don't have time or need to review APM logs, Supportability Metrics, or the Status Tab.

New Myth: I will make it part of my Administration responsibilities to review each of the above to get a better understanding on what is typical behavior and be closer to problems when they first occur.

Central Myth #6: APM CE can monitor as many transaction as as I want.

New Myth: I will work with my application owners to only monitor what is really needed so to avoid filling up the APM database with unnecessary data and produce unneeded reports and metrics.

Central Myth #7: Using a virtual environment for APM means hands-free administration

New Myth: I will monitor the cluster health of my APM virtual environments with more scrutiny than my physical environments. They are more likely to hit limits faster and experience more performance problems,

Central Myth #8: There is no reason to read the documentation, release notes, KBs, community posts, attend APM Office hours etc. Whatever is out there is hard to find, unclear, and incomplete.

New Myth: Look at the resources as a first resort. APM Tech Docs and Support work hard to make documentation, release, and KBs as clear, complete, and current as possible. If this is not the case, then post a comment on wiki.ca.com with the details or send me a note and I will see it gets rectified. APM Office Hours is your time to directly ask APM Product Management and Support your questions and concerns

Central Myth #9: Reactive is cheaper and less time-consuming than proactive.

New Myth: Proactive administration stabilizes environments and anticipates future application monitoring needs.

Central Myth #10: I am going to be here for some time. There is no need to document operating tasks and best practices.

New Myth: Circumstances are always changing. Keep your run book, APM monitoring strategy/roadmap, and transaction definition "spreadsheet" current.

There are many more central myths that can be added here. Perhaps in a future article.

Reference:

Molloy, Michael Experiencing the World's Religions: Sixth Edition. New York: McGraw-Hill Education 2012

Discussion Questions:

1. What are your APM Administration Central Myths?
2. Have these been ever challenged? What happened as a result?
3. How frequently should one evaluate their APM Administration central myths?

3.11 Successfully Configuring and Troubleshooting RTTM with User Groups

URL: <https://communities.ca.com/message/241840970#241840970>

Setting up and Troubleshooting RTTM and User Groups

APM 9.7 added a long-requested functionality -- support reporting on User Groups under RTTM. This is a powerful capability that will help determine quickly which collection of users are having issues with specific transitions. But in dealing with a customer case and reading the documentation, I realized that more needed to be documented in this area.

Two KBs were created which are listed below. This article provides details than what is listed therein.

Here is the essence of what I learned in the process:

Lesson #1: It helps to have a User Group Identifier such as Client IP. After doing this, the customer started to see user groups.

Lesson #2: User Group names on the Investigator and APM CE (CEM) GUI may not match. TIM User groups in the Investigator will have a *tim-* prefix while in the CEM MOM GUI, they will not.

Lesson #3: The alias and the pattern in the RTTM Matcher do not need to match.

It was unclear from reading the documentation if the alias was used in matching or not. But it is basically a description and nothing more.

Lesson #4: Troubleshooting changes depending if CEM is doing user grouping with or without RTTM under User Grouping enabled.

Lesson #5: The key to troubleshooting is examining two artifacts.

These are the following two files on the TIM:

1) btstatsconfig.xml in TIM_Home/Config

```
<?xml version="1.0" encoding="UTF-8"?>  
  
  -<BTStatsConfig reportUserGroup="1" reportUrlPath="0"  
maxInstPerInterval="600" enabled="1">
```

In our case, we found out that this changed value was not being written out to this file.

It must have reportUserGroup enabled (i.e set to 1). Note that reportUrlPath may or may not be set

2) btstats.xml files in the TIMs "data/out/btstats" directory

Since these files are processed quickly, one have to stop the TIM, run a "watch cp" command to copy the file, and the like. The important thing is to verify is if user group information is being passed in the RTTM file. There will be a stanza line like below if this is the case.

```
<BtInstance count="integer" userGroup="group name"... />
```

Lesson #6: Sometimes deleting the User Group from the APM CE GUI helped.

This allowed a new user group to be created. And the RTTM data files populated correctly.

I hope this is helpful. Please add any additional tips and I will get these added in the appropriate documentation

References:

1. <http://www.ca.com/us/support/ca-support-online/product-content/knowledgebase-articles/TEC1749361.aspx> -- I have read the "Configure User Groups under RTTM" documentation. But it is not working as expected.
2. <http://www.ca.com/us/support/ca-support-online/product-content/knowledgebase-articles/TEC1602488.aspx> -- RTTM under User Groups Troubleshooting

Questions for Discussion:

1. Are you using the RTTM User Group functionality?
2. How easy was it to setup? Did you do some of the above in resolving your setup issues?
3. What else should I discuss?

3.12 TIM Trace Options -- The Missing Pages Part 1

TIM Trace Options -- The Missing Pages Part 1

Introduction

Proper use of the TIM trace options can be useful in solving a variety of issues. Some of these include:

- 1) If seeing network traffic at all
- 2) Which servers the TIM is seeing traffic from.
- 3) Is the traffic to the TIM two-way showing the HTTP request and response?
- 4) SSL issues such as ciphersuites and TLS versions used.
- 5) Network data quality issues (OOO packets, packet resets/terminations, etc.)
- 6) Transaction definition issues such as transaction counts, response time, transaction component issues, etc.
- 7) HTTP Plugin issues.

This article gives an overview on TIM Trace options and responsible use.

Overview

The TIM Trace options are configured on each individual TIM. This is done by going to the **System Setup page**, then select TIM, **Configure TIM Trace Options**.

The appropriate programs read the trace options checked and saves it in a text file named tim.options in the TIM config directory. This file cannot be edited. The TIM web server is notified of the changed trace options

Here are some guidelines on TIM trace options.

1. TIM Trace options should be turned off unless needed for debugging. This is especially true for busy production networks. If left on, the TIM may restart due to excess CPU/Memory. On MTPs, it may also result in a full RAMdisk directory. TIM logging levels should be checked early if there are TIM performance issues.

2. If debugging with Trace options, only turn on the trace options truly needed. This is discussed in Part 2 to be published next month.

3. Note that certain TIM settings may add or subtract what appears in the TIM log or the content of the messages. Especially of note are the following:

- Enabling plugins

FlexPlugin/Enabled = 1

FlexAMFParser/Enabled = 1

FlexAMFXParser/Enabled = 1,

and XmlPlugin/Enabled = 1

- Enabling tracing

SiteMinder/Trace = 1

AllXmlParams = 1 if wanting every XML element and attribute rather than those that match a parameter definition

BtStats/LogLevel = 1 or 2

(1 = Messages are written to the TIM log when BtStats files are written,

or 2 = Same as messages, verbose mode)

- Autogen/Trace = 1 or 2. This is for transaction discovery. 1 = Traces autogen definitions creation 2 = Additional detail messages

3. Use the TIM Trace Filters to limit traffic to a particular client/server IP address, browser language setting, or HTTP parameter.

4. Try to make TIM trace options the same across all TIMs. This produces consistent logging.

Additional Information.

Here is an incomplete where to look to find more:

- <https://docops.ca.com/pages/viewpage.action?pageId=275682398> -- Using TIM Trace options for troubleshooting.

- <https://docops.ca.com/display/APMDEVOPS101/Transaction+Identification> -- Using TIM Trace options for transaction identification.

-

<https://docops.ca.com/display/APMDEVOPS101/Setting+up+CEM+to+support+nCipher>
-- Using TIM trace options with nCipher.

- <http://www.ca.com/us/support/ca-support-online/product-content/knowledgebase-articles/TEC1281821.aspx>

-- Which TIM trace options to use when debugging issues. We will return to this next month.

- <http://www.ca.com/us/support/ca-support-online/product-content/knowledgebase-articles/tec610515.aspx>

- How to parse a TIM log by IP address.
- <http://www.ca.com/us/support/ca-support-online/product-content/knowledgebase-articles/tec600884.aspx>

- TIM loggings can create a MTP no space or TIM restart condition.

