

# Expanding Your System past just a PI Historian – A 2016 Update

Bruce McCamant, TSI  
September 15, 2016



# Triencon Services, Inc.



An Energy Services Company...

- Providing services and consulting

- Average experience in power generation, transmission/distribution, market ops/trading - 25+ years

- 20+ years experience with PI

We all started as customers who installed PI and served as PI administrators

Experienced in integrating PI with a wide variety of EMS, SCADA and DCS platforms

Other services include—

- IT Consulting and services

- Virtualization technologies – Hyper-V, Vmware, Vblock

- NERC CIP – consulting, remediation, integration

Microsoft Gold Partner, Dell Preferred Partner

# Note to Audience

This presentation is intended to meet you where you are at currently – a drive-by if you will, not in-depth by any means

- You may be
  - An experienced PI admin who is already at PI 2016
    - Feel free to (gently) correct / share
  - A PI admin who has yet to upgrade to 2015 or 2016 – or even 2012
    - It might offer a glimpse of things to consider that may be applicable
  - Only a PI user
    - It might give you insight to what's under the hood, an insight to other features/tools you might want to consider for your company
  - A potential PI customer considering PI
    - Varnished opinions not given here. I point out the real stories – of a platform that is still by far the leader in the industry

## Last year...

- We talked about PI 2015
- We discussed all the features that came with it that you should be using
  - And all of the ‘overhead’ it takes when you do an HA deployment and expanded use of AF and Coresight
- In particular we discussed
  - High Availability configurations
  - Windows Authentication

## Since then...

PI 2016 has been released – of course!

When we say PI that includes

- Archive Server
- AF Server
- AF Client
- Analysis
- Client/Visualization Tools
  - Datalink
  - Coresight

There may be a clear and perhaps overriding factor to be considered for upgrading to 2016 – Security Update to PI AF

Self-reported vulnerability in May of 2016 –

<https://techsupport.osisoft.com/Troubleshooting/Alerts/AL00301>

But more importantly known to ICS-CERT as ICSA-16-166-02 (<https://ics-cert.us-cert.gov/advisories/ICSA-16-166-02>)

If your company follows CIP guidelines in a best practice model in non-CIP spaces, this is identified as a medium grade vulnerability by CERT scoring

The vulnerabilities deal primarily with access to port 5459 used by PI OLEDB Enterprise; PI Coresight 2013 and earlier; PI Webparts 2010 and earlier

Check your configuration and consult the alert for possible mitigation options

# PI 2016 Features

Along with fixes, 2016 has new features

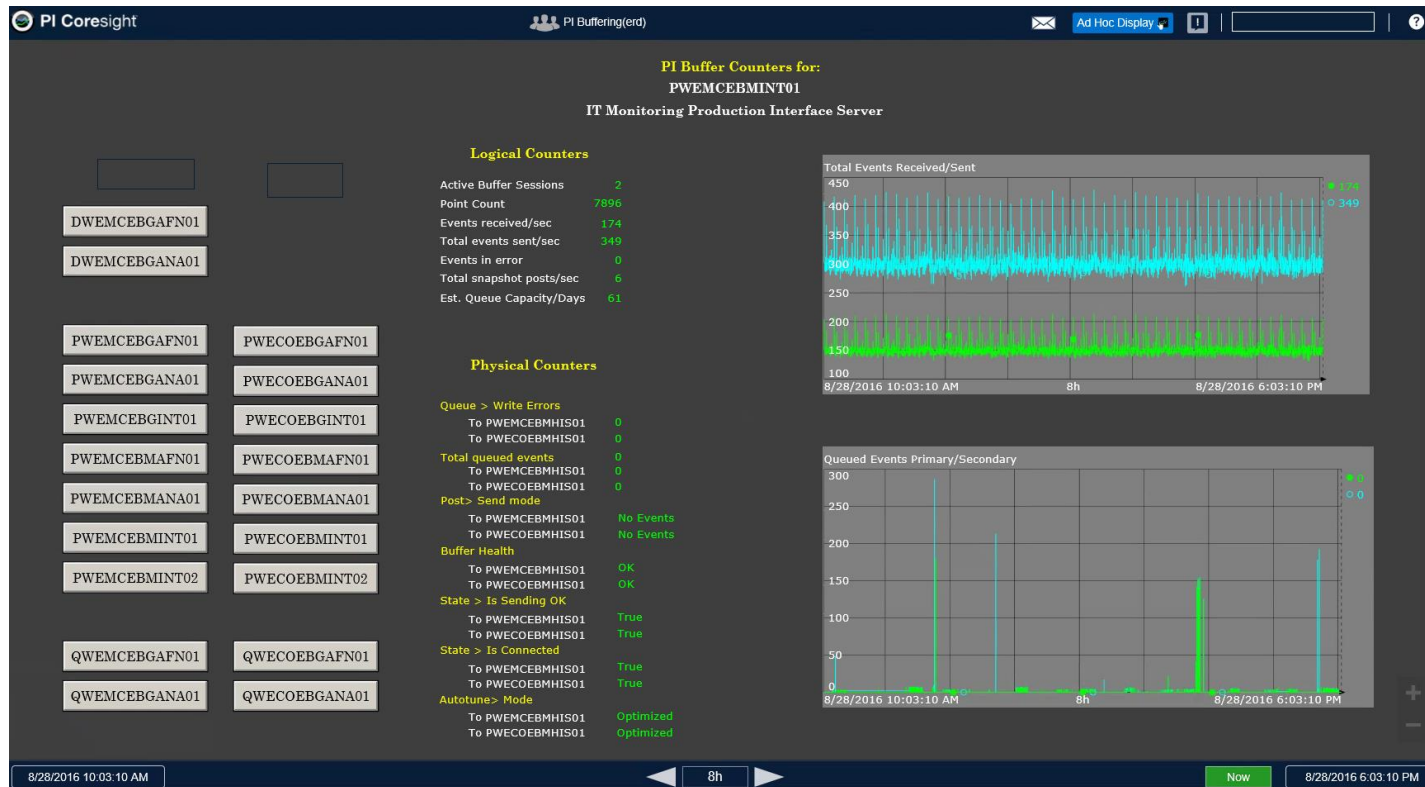
- **Archive**
  - Archive server self-detects corrupt archives; users can reprocess them online for repair from PI-SMT
  - Improvements in certain methods of reading and writing data for tags
- **AF/Analysis**
  - Sets the stage for the next version of PI Notifications
  - Attributes now can have traits. Traits can include limits
  - New functionality to retrieve and act on data in archives
  - New functions in Analytics library
  - Event Frames can have multiple start triggers

# PI Analytics

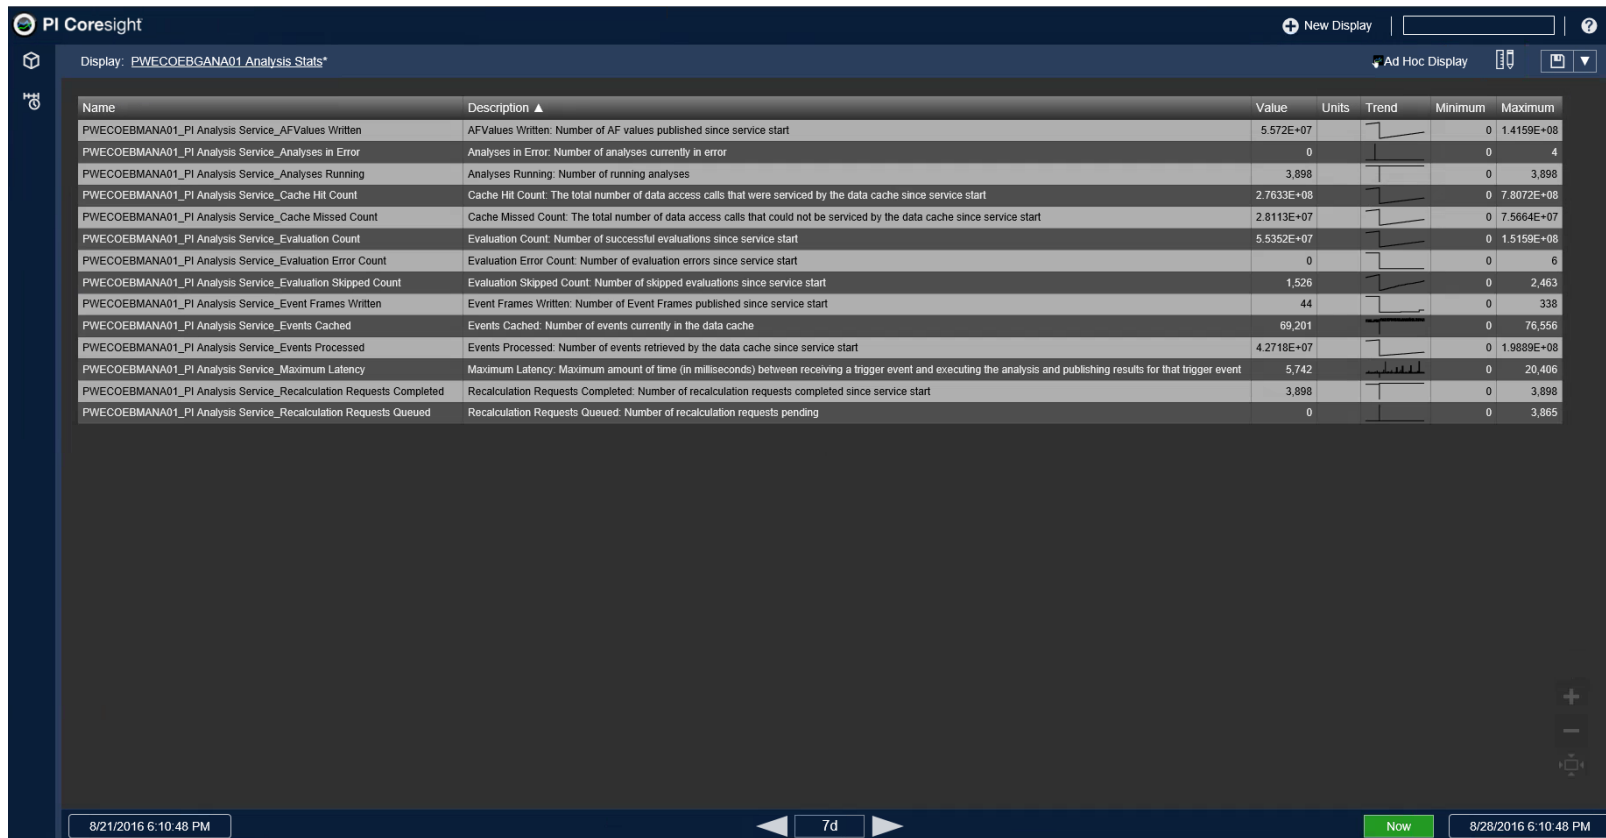
- Note that if you have intend to use PI Analytics to write values to tags – and you have a PI archive collective – you need to be using PI Buffer Subsystem
- PI Analytics just like an interface in this way, not like PI Performance Equations
  - Analysis does not run on the Archive server
  - Without buffering, results will only be written to the primary archive server; data needs to be replicated and buffered
- PI Buffer Subsystem and PI Analysis both come with their own set of Perfmon tags to monitor health – consider a Perfmon instance to bring back some of this data to help you monitor Analytics health



# Monitoring of PI Buffering



# Monitoring of PI Analytics

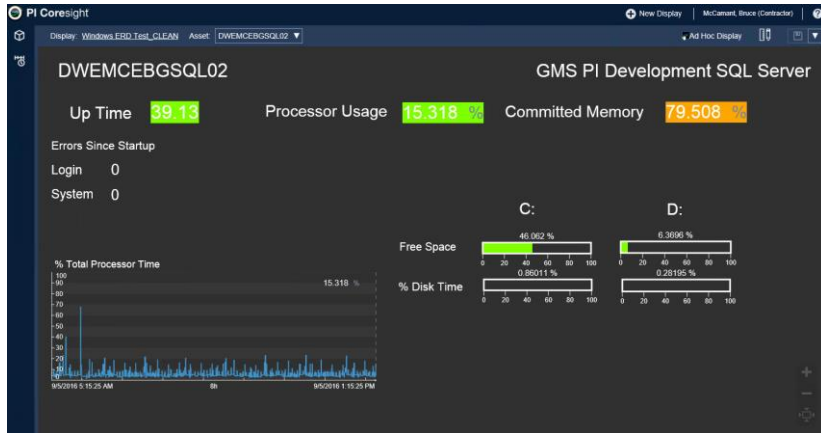


The displays on the prior slides were taken from PI Coresight 2016

- 2016 is a major new release of the product
  - Think ‘web-based ProcessBook’ - but... it is the first release
  - It does have its ‘warts’
    - ‘Nagging’ things that are crucial to ProcessBook – don’t work the same in Coresight
      - The background color in Coresight is not in the color palette for symbols
      - Can’t sync sizes of multiple symbols
      - Gauges – somewhat limited in capability of formatting compared to ProcessBook
      - The ability to set different time ranges on multiple trend symbols that was in Coresight 2015 – not available in 2016

# PI Coresight 2016

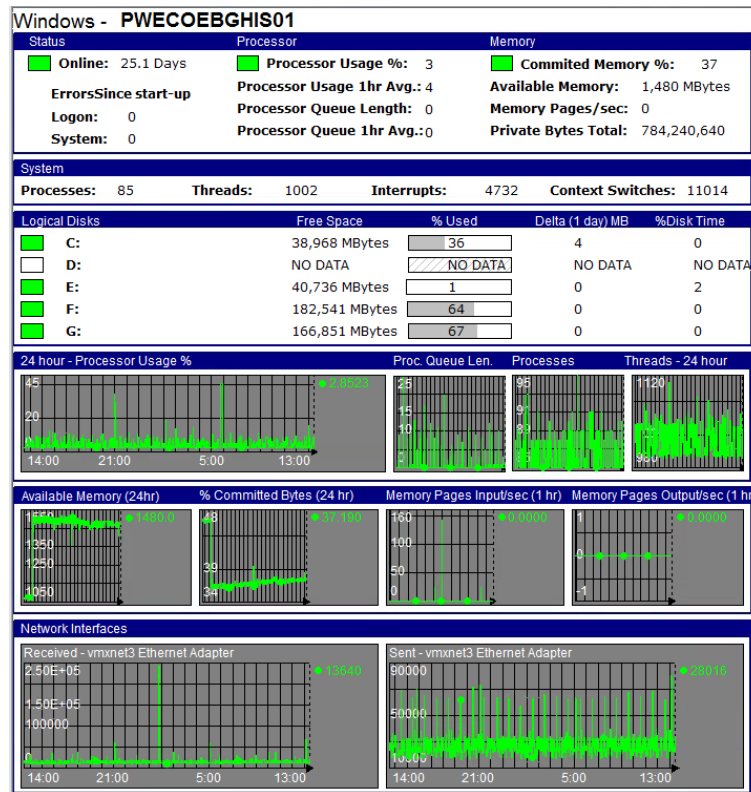
## Background Color



# PI Coresight 2016 — Background Color Adventure

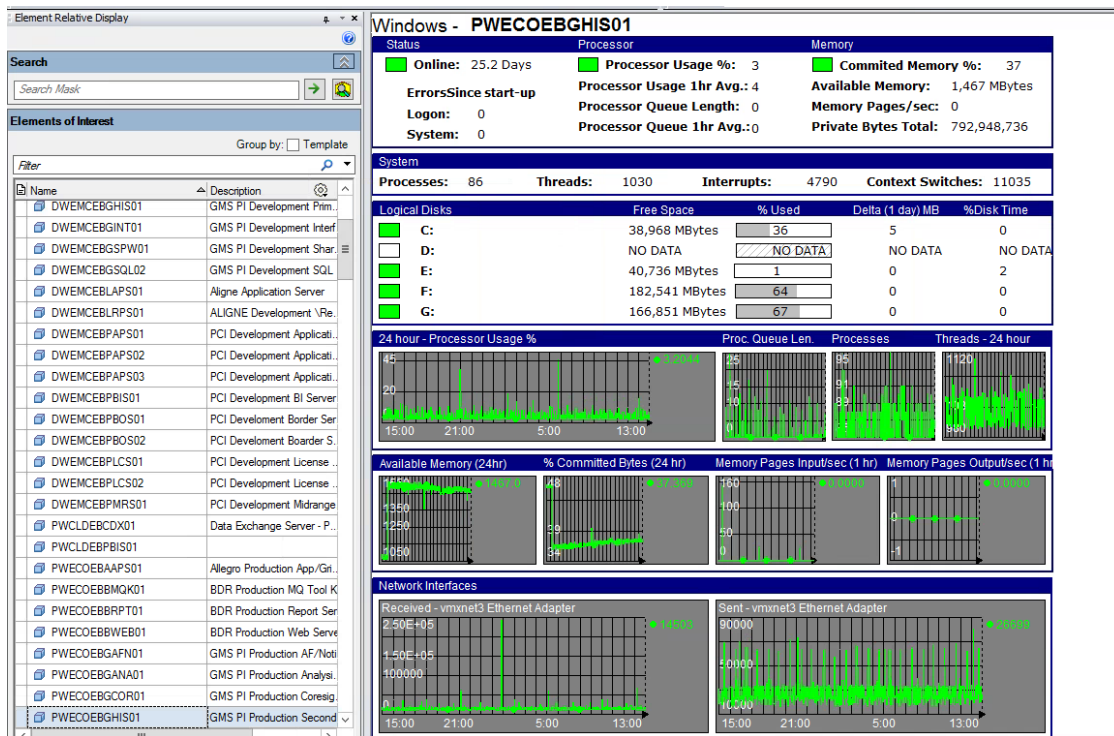
This is also an example of where AF can take you –

Here is an example of a ProcessBook display – which previously was built with a display file for each asset



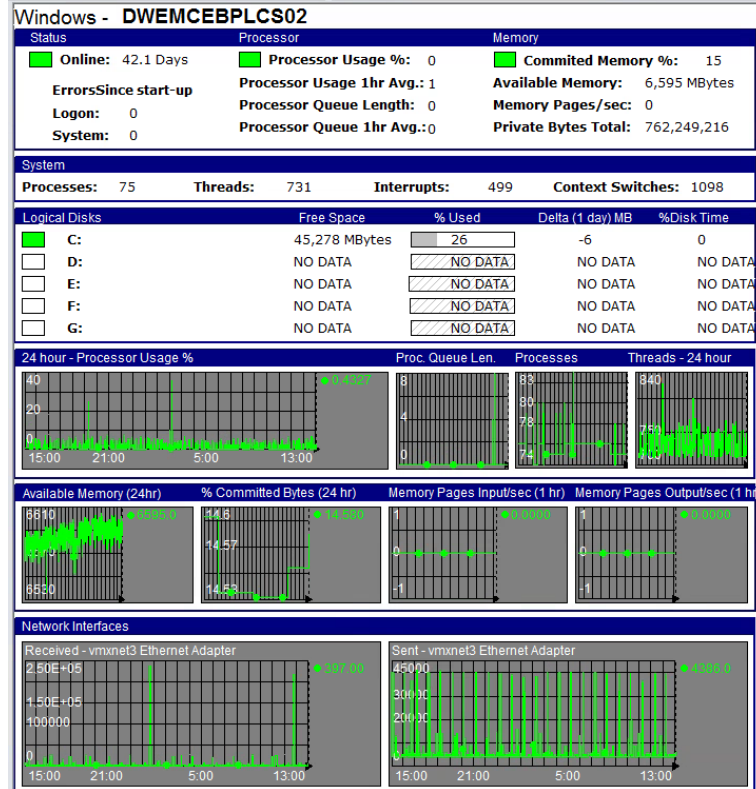
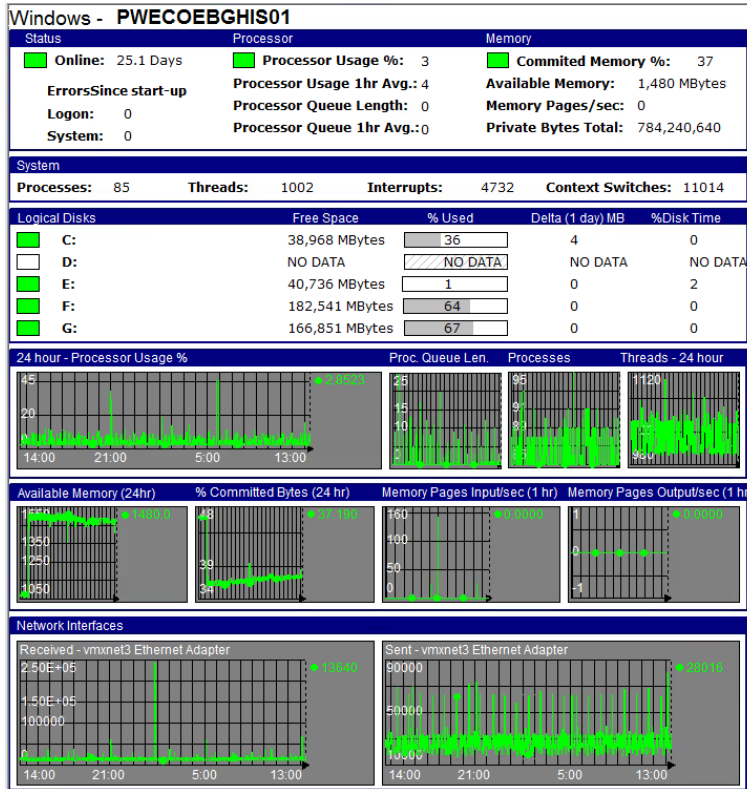
# PI Coresight 2016 — Background Color Adventure

We take this display and convert the display objects to point to an Element Relative model in AF



# PI Coresight 2016 – Background Color Adventure

This is nice – but  
not every asset  
has the same  
configuration...



# PI Coresight 2016 — Background Color Adventure

But we were forward-thinking in our AF model – used the Excluded property in our model

PWECOEBMHIS01

General Child Elements Attributes Ports Analyses Version

Filter

Name	Value	Time Stamp
LogicalDisk(C:)_Free Megabytes	39011 MB	9/5/2016 1:24:15 PM
LogicalDisk(D:)_Low Limit	10 %	1/1/1970 12:00:00 AM
LogicalDisk(D:)_Critical Limit	5 %	1/1/1970 12:00:00 AM
LogicalDisk(D:)_Eval Time	5	1/1/1970 12:00:00 AM
LogicalDisk(D:)_Score	Excluded	1/1/1970 12:00:00 AM
LogicalDisk(D:)_% Disk Time	Excluded	1/1/1970 12:00:00 AM
LogicalDisk(D:)_% Free Space	Excluded	1/1/1970 12:00:00 AM
LogicalDisk(D:)_%Disk Space Used	Excluded	1/1/1970 12:00:00 AM
LogicalDisk(D:)_Avg. Disk Queue Length	Excluded	1/1/1970 12:00:00 AM
LogicalDisk(D:)_Current Disk Queue Length	Excluded	1/1/1970 12:00:00 AM
LogicalDisk(D:)_Disk Reads/sec	Excluded	1/1/1970 12:00:00 AM
LogicalDisk(D:)_Disk Writes/sec	Excluded	1/1/1970 12:00:00 AM
LogicalDisk(D:)_DiskSpaceDelta	Excluded	1/1/1970 12:00:00 AM
LogicalDisk(D:)_Free Megabytes	Excluded	1/1/1970 12:00:00 AM
LogicalDisk(E:)_Low Limit	10 %	1/1/1970 12:00:00 AM
LogicalDisk(E:)_Critical Limit	5 %	1/1/1970 12:00:00 AM
LogicalDisk(E:)_Eval Time	5	1/1/1970 12:00:00 AM
LogicalDisk(E:)_Score	100	9/5/2016 1:25:45 PM
LogicalDisk(E:)_% Disk Time	3.450903 %	9/5/2016 1:21:05 PM
LogicalDisk(E:)_% Free Space	99.4628372192383 %	9/5/2016 1:25:45 PM
LogicalDisk(E:)_%Disk Space Used	0.537162780761719 %	9/5/2016 1:25:45 PM
LogicalDisk(E:)_Avg. Disk Queue Length	0.0336790829896927 %	9/5/2016 1:25:45 PM
LogicalDisk(E:)_Current Disk Queue Length	0 %	9/5/2016 1:24:25 PM

Group by: ☐ Category ☐ Template

Name: LogicalDisk(D:)\_% Disk Time

Description:

Properties: Excluded

Categories: Disk Info

Default UOM: percent

Value Type: Single

Value: Excluded

Data Reference: <None>

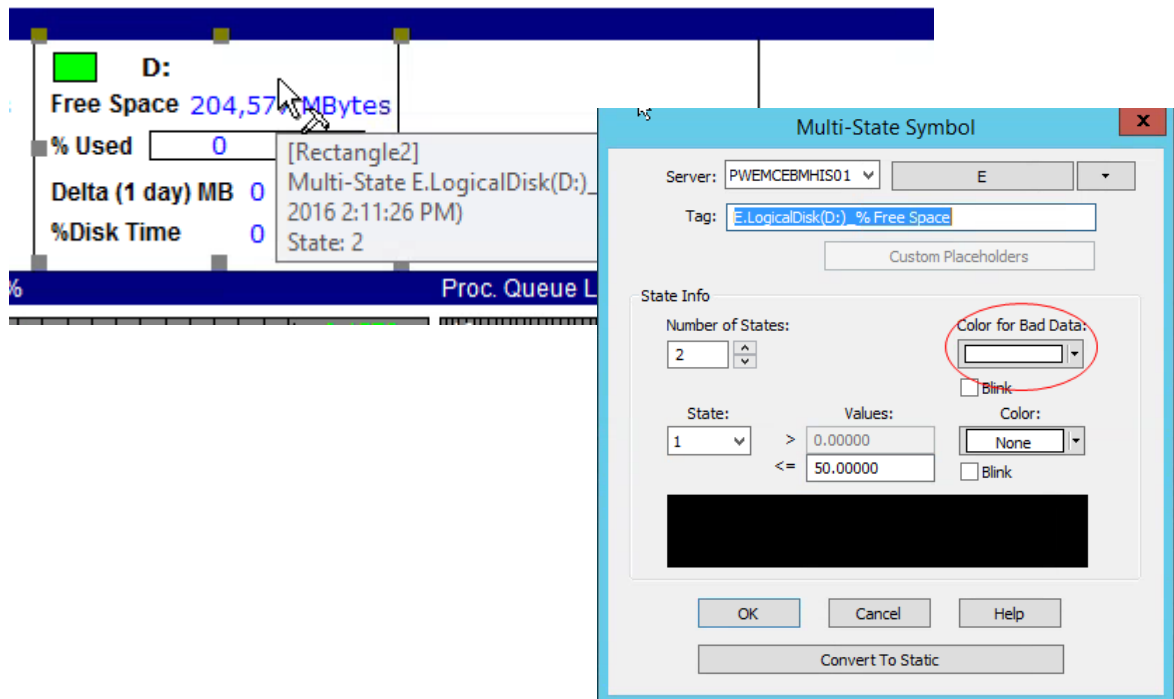
Settings...

Limits Forecasts



# PI Coresight 2016 — Background Color Adventure

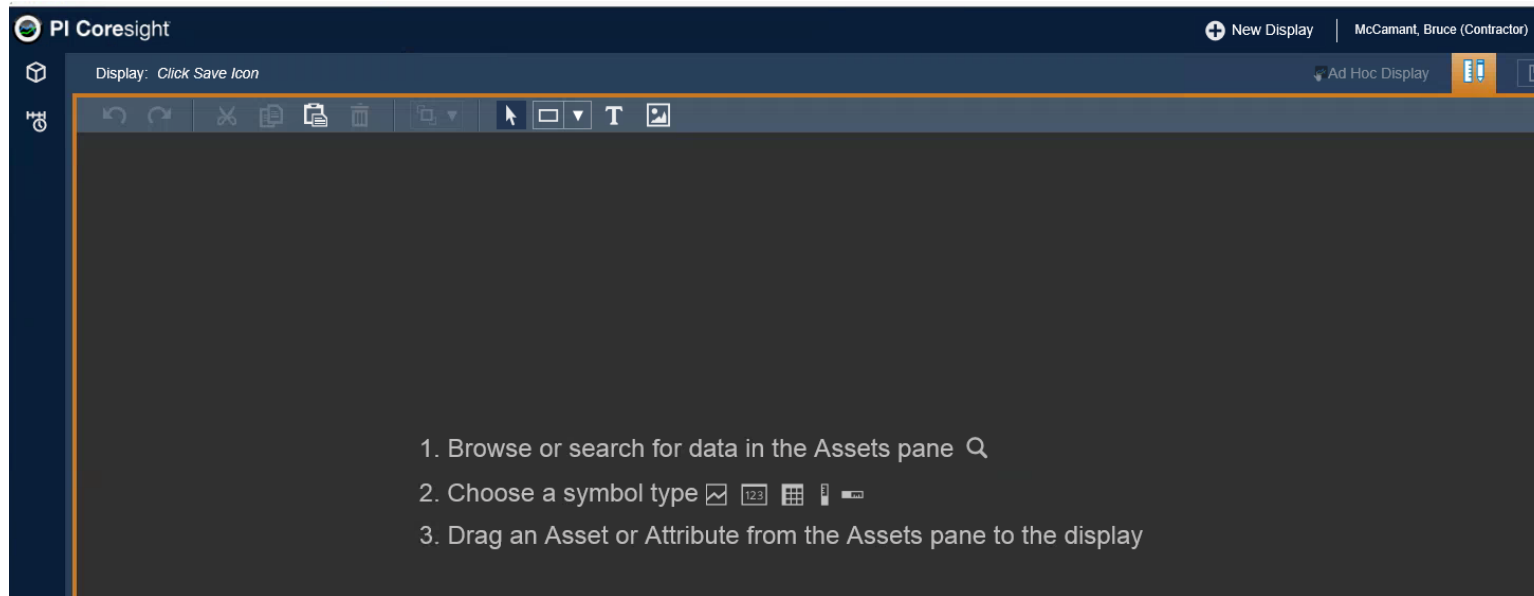
We use a multi-state on a rectangle object to hide 'Excluded' attributes – resulting in a cleaner display



# PI Coresight 2016 — Background Color Adventure

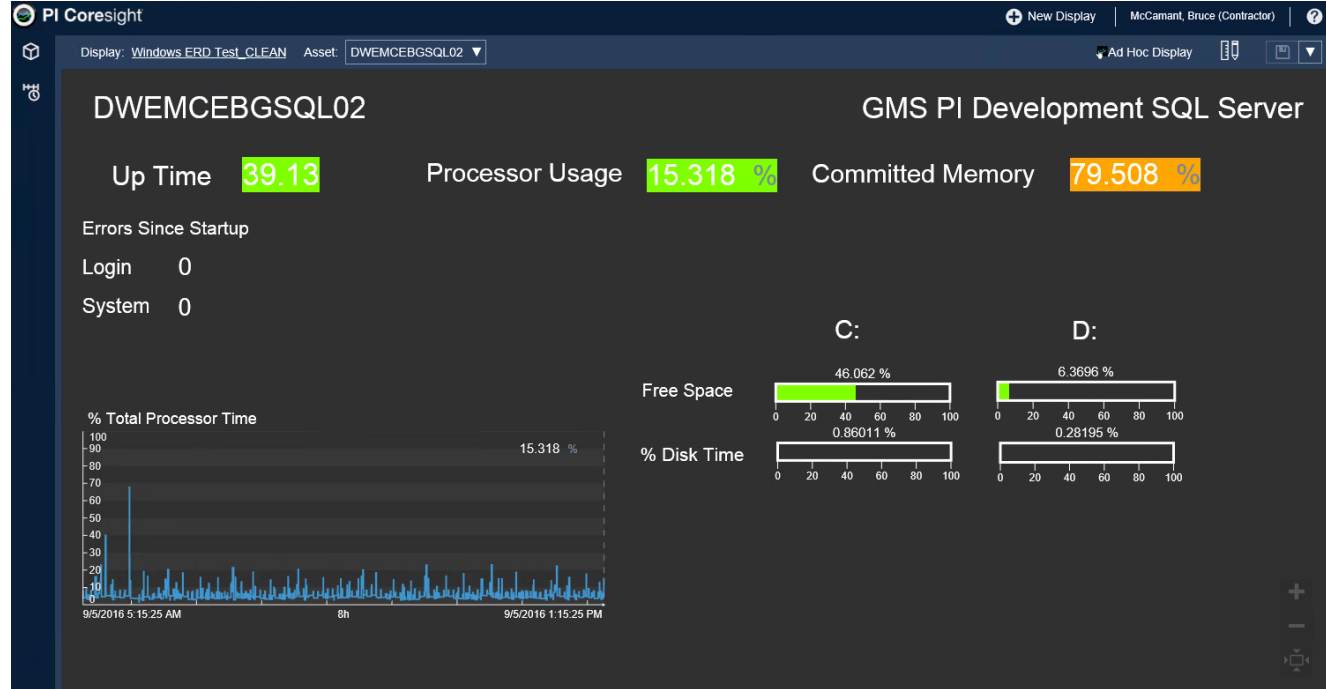
So now I have Coresight 2016!

- I have drawing objects
- I have multi-states



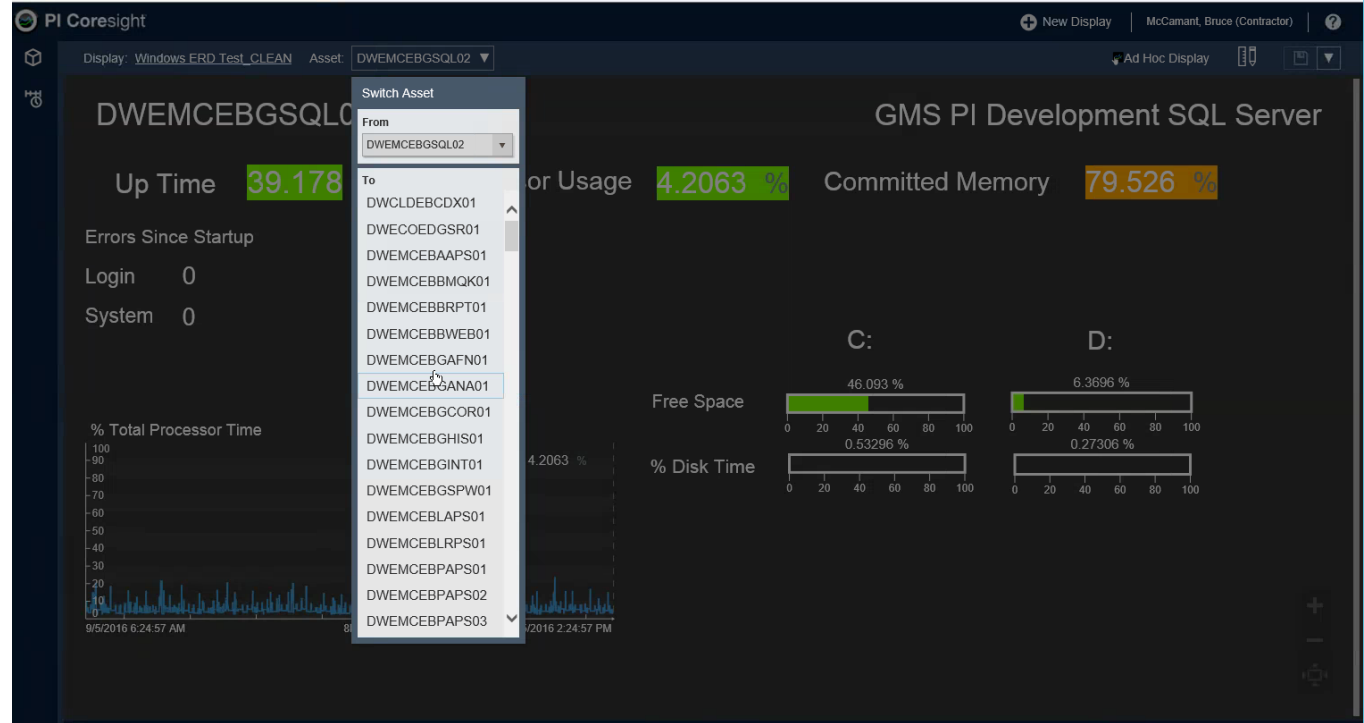
# PI Coresight 2016 — Background Color Adventure

I am to going build the  
equivalent of my element-  
relative display in  
Coresight



# PI Coresight 2016 — Background Color Adventure

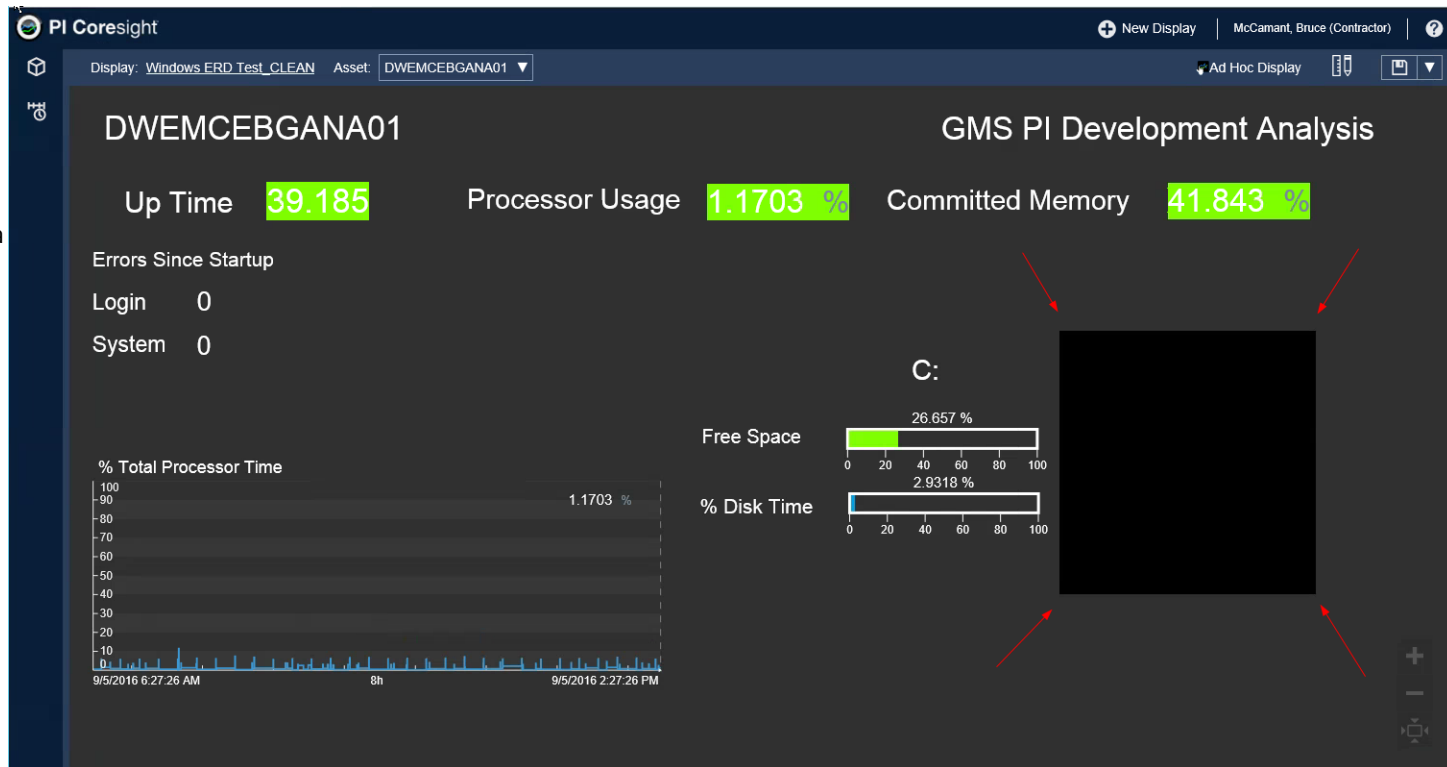

I have an Asset  
Picker – that comes  
alive based on  
content – not even a  
separate add-in!



# PI Coresight 2016 — Background Color Adventure

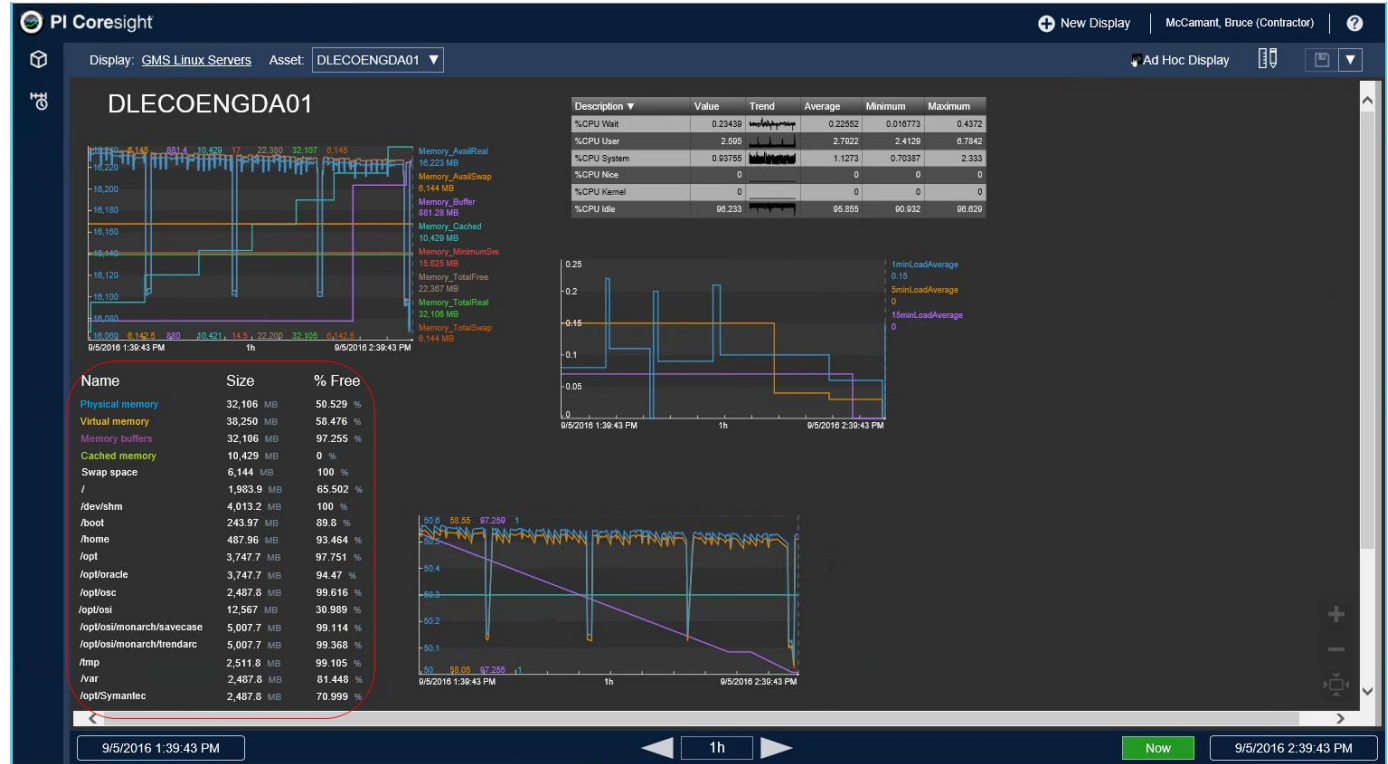
Cue the 'thud' sound here...

This has been reported to Tech Support – and listed as an issue on the UserVoice forum – you can actually reach from a Coresight display



# PI Coresight 2016 – Sizing of Objects

Consider the ‘table’  
at the bottom left –  
actually composed of  
individual value  
objects



# PI Coresight 2016 – Sizing of Objects

- These handles for these symbols – limited functionality
- When you grab the handle, the cursor indicates you are controlling in one dimension – when in reality you are changing size but maintaining aspect ratio
- While you can select multiple symbols – sizing is not an option you have
- If you select one, it is very difficult to determine size of one that you can apply to others – no sizing feedback available.
- Reported on UserVoice as well

Name	Size	% Free
Physical memory	32,106 MB	50.527 %
Virtual memory	38,250 MB	58.474 %
Memory buffers	32,106 MB	97.255 %
Cached memory	10,429 MB	0 %
Swap space	6,144 MB	100 %
/	1,983.9 MB	65.502 %
/dev/shm	4,013.2 MB	100 %
/boot	243.97 MB	89.8 %



# PI Coresight 2016 – Justified

All things considered – it is still a very valuable tool that definitely should be part of your data visualization strategy

- Very stable
- You can publish ProcessBook displays here so that you don't have to give everyone a ProcessBook install
  - Eliminates upkeep of software on individual user workstations – no PI software needed
- Very easy to use for users with minimal training compared to ProcessBook
- You can control content to users by
  - Control the content they can access via AF security
  - You can provide 'limited' users access via URLs that limit what they can do – include arguments to control access to time bars, use kiosk mode, etc.



# PI and Authentication – an Update

When we spoke last in 2015 we discussed the role of Kerberos authentication in PI

- Let's step back and talk about authentication methods in PI – for the Archive server
  - PI user account – the most basic
    - Should not be in use – see self-reported <https://techsupport.osisoft.com/Troubleshooting/Alerts/AL00206>; This was also disclosed by NERC in 2009 - <http://www.nerc.com/pa/rrm/bpsa/Alerts%20DL/2009%20Advisories/A-2009-10-02-01.pdf>
  - PI trusts – better, but still not the most secure
    - Used typically where domain authentication is not available
      - Should be strengthened by being more detailed
        - » Specify allowed applications – can also include domain account
      - In a control center or 'critical' use scenario – consider it as a backup for if and when Active Directory services become unavailable for an extended period – limited audience
        - » Using visualization against tags only – no AF
  - Windows authentication against Active Directory is most secure

# PI and Authentication – an Update

Let's review the products and see what authentication is available

- PI Archive – ~~PI user~~, PI Trust, Windows Authentication
- PI AF – Windows Authentication\*
- PI Analysis – Windows Authentication
- PI Coresight – Windows Authentication\*
- PI WebAPI, PI Crawler – Windows Authentication

# PI and Authentication – an Update

## Singling out PI AF and PI Coresight

- AF depends on AF client installed and running on the node you are logged on and using
  - It hands off user access to existing Archive data content to the authentication method configured – which could be a Trust
  - In this case – archive data can be viewed that uses the trust
  - If you want to develop any new content that has PI Point references as attributes – the trust will not be sufficient for access because AF does not use trusts

# PI and Authentication – an Update

## Singling out PI AF and PI Coresight

- It is a possible scenario whereby a trust to the Coresight server can be used for data access
  - Again – it would only work on existing content (a ProcessBook display for example)
  - And again - if you want to develop any new content in Coresight the trust will not be sufficient for access because Coresight does not use trusts

# PI and Windows Authentication using Kerberos

First – a ‘high’ level view of Kerberos

- Involves secret keys and tickets to allow secure authentication in a ‘less than secure’ network – part of Active Directory but a separate application
- Users/Servers/Applications can prove their identity to one another
- Kerberos enables ‘double-hop authentication’ – an app using it can authenticate on behalf of a user to another application
  - For PI
    - AF gets a user access to PI archive data
    - Coresight gets a user access to both PI archive and AF data

# Kerberos Setup for PI

- READ THE MANUALS FIRST!!!
  - OSIsoft has done an excellent job of laying out what needs to be done. If you need to, copy pertinent sections out and email to your domain admin doing the AD configuration
    - There is a good chance – it will be their first time
- If you are running a standalone environment you are using the default local accounts the apps install
- If you are running an HA environment you are then using a domain service account to run AF and/or Coresight as opposed to local system.
- For Kerberos setup for use with PI apps we are required to –
  - Set up a name for the service that Kerberos can refer to for authentication
  - Set up the ability for Kerberos delegation for each server – can do it generally or by server and app

# Kerberos Setup for PI

First – create/verify Service Principal Name(s) (SPN) for the applications; basically a handle that Kerberos can know the service by for authentication – affiliated with the application server(s) and the service account being used

- We actually enable the account used to both read and write the SPN for the application when it starts it up
  - For PI AF we set this for the AF application
  - For PI Coresight we basically do this for HTTP access – not really a Coresight ‘app’
- If app is standalone using local accounts, SPNs should exist upon creation – but should be verified
- For HA using domain accounts – a domain admin runs commands to create them in AD

# Kerberos Setup for PI

Next – the ability to perform Kerberos delegation by the server(s) must be enabled – must be done whether standalone or HA

- This can be done one of two ways
  - Give the server the ability to do all things Kerberos
  - Or in specifics – name the servers and the apps
    - Example – for Coresight you would list each server and app
      - AF servers and the service for AF – by the SPN handle
      - Archive servers and the SPN for PI archive (self-registered since we run PI archive as local system)



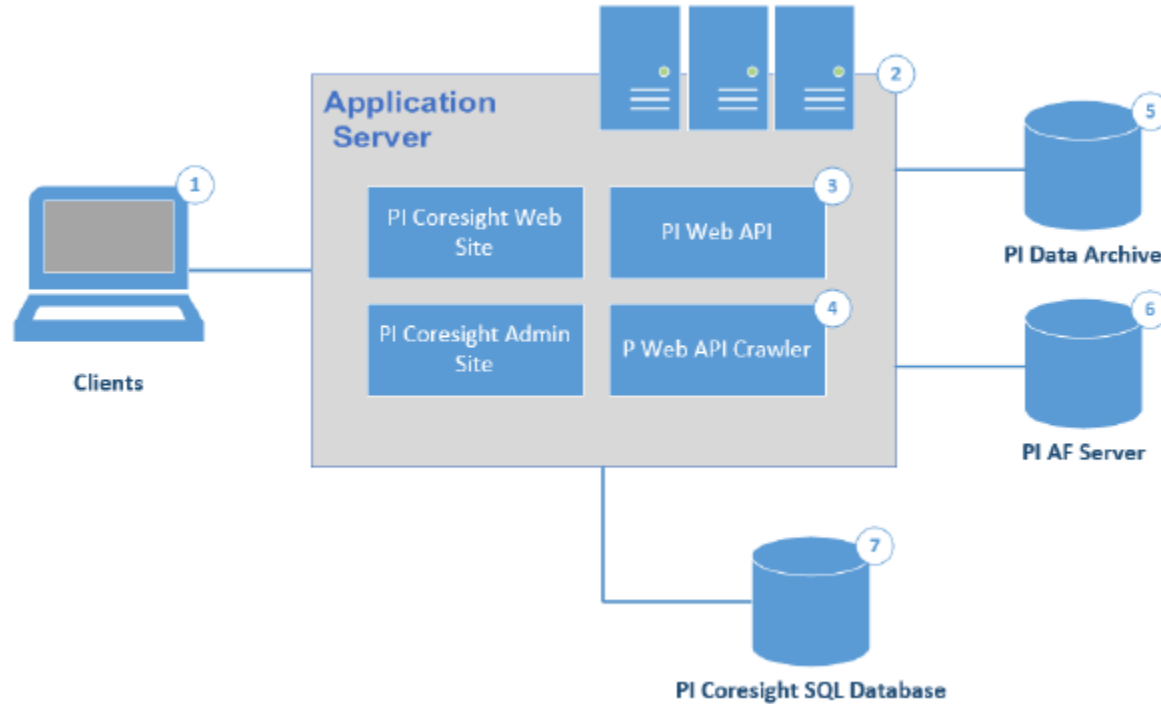
# PI and Windows Authentication using Kerberos

Key points –

- For PI AF –
  - Option exists to not use Kerberos; This is done in AF configuration file. Good news – if you configure to have AF use Kerberos, it is possible for AF to try NTLM if Kerberos fails. (Default setting is to use Kerberos)
  - If you intend to use an AF Linked Table to the SQL Server, you will need to verify/setup that the domain account used to run the SQL Server Engine application is set up for SPN privileges as well;
    - SQL Server in an HA arrangement (mirror, Always On) may well need it as well just for SQL to work HA; just want to follow up if AF Linked tables are in your plans
- For PI Coresight – If Kerberos is configured, there is not a really a means to fallback to NTLM.

# PI Coresight

From the  
Installation and  
Administration  
Guide



# PI Coresight

## (1) Clients

Clients access via browser options and platforms – desktop and mobile

(2) PI Coresight application server  
The application server handles all application operations between users (clients) and PI Data Archive servers, PI AF servers, Microsoft SQL Server, and PI Web API.

## (3) PI Web API

The PI Web API provides cross-platform, multi-user programmatic access to PI System data.

## (4) PI Web API Crawler

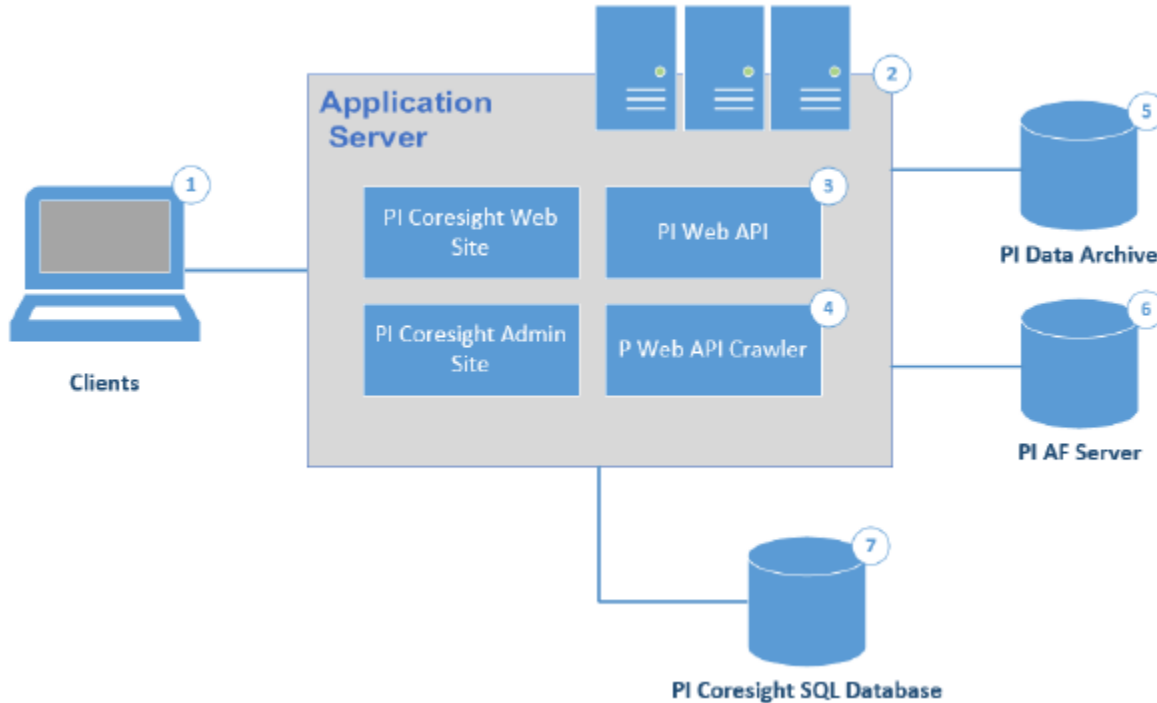
The PI Web API Crawler enables the search functionality of PI Coresight by periodically crawling the registered servers for metadata and creating a search index.

## (5) PI Data Archive server

## (6) PI AF server

(7) PI Coresight database in Microsoft SQL Server

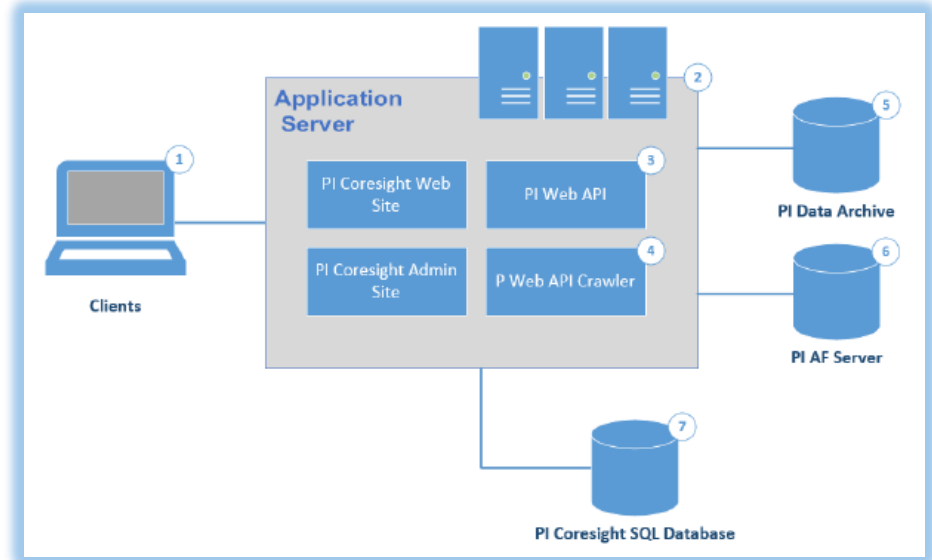
PI Coresight stores user display settings and definitions in a Microsoft SQL database. Display definitions include such data as display name and display owner, symbols on the display, user permissions, etc.



# PI Coresight – new display

When a user searches for data in a new display –

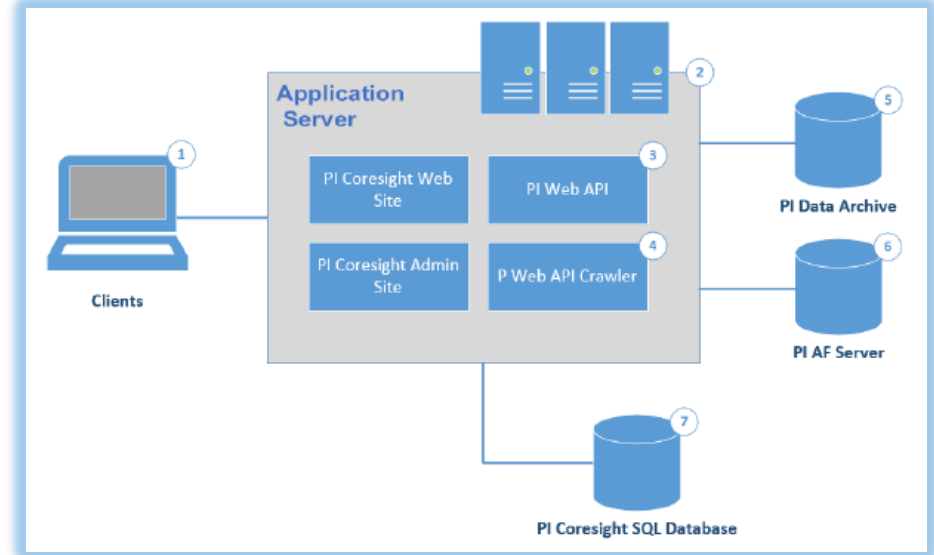
1. The request from the user is handed off to the PI Web API, which populates results from the local index created by the PI Web API Crawler
2. The user after selecting data chooses which symbol to create for the data item(s) on the display
3. The creation of the symbol triggers the request for PI data to the source – PI Archive or PI AF. This is where the application would use Windows Integrated Security (WIS)/Kerberos to present the client's credentials. If successful and allowed, data is returned for the symbol(s)
4. When the user saves the display, the display definition settings are saved in the PI Coresight database in Microsoft SQL Server.



# PI Coresight – existing display

When a user opens and existing display –

1. The Coresight application sends a request for the display definition from the Microsoft SQL Server Coresight database
2. When the symbols are populated, that triggers the request for data from PI Archive or PI AF – which authenticate the client and their access



# PI Coresight – Authentication

Notice I did not mention Kerberos for existing display...

I mention this because some of you may see Coresight 'work' if you are still using trusts – may be confusing

- For those of you inclined to not use any of type Windows authentication – not possible
  - There is a local group on the Coresight Server called 'Coresight Users'. Windows credentials only – PI trusts are not recognized
- It is possible for data access to the Archive server to happen by way of PI Trusts.
  - You could have a trust to the Coresight server and a trust for a client workstation. If they are in the Coresight Users group this client by virtue of the trusts could see data on existing displays
  - This method takes a lot of care and feeding on your part – and WIS is much more secure than PI trusts

# PI Coresight – Authentication

## Bottom Line

- You have to use WIS to grant user access for Coresight to begin with
- Without it – You (or any other users) cannot build content
  - The PI Web API and Search Crawler are part of the application – not optional

# PI Web API – A Recommendation

- Many of you may have other applications collecting data from PI
- Many of you still use what would be considered ‘legacy’ methods against the Archive Server
  - PI SQL
  - PI OLEDB Provider
  - PI JDBC Provider
- These specifically utilize the PI-SDK – which OSIsoft keeps telling us – will be deprecated (no major enhancements)



## PI Web API – A Recommendation

- The PI AF-SDK is already proven to retrieve archive data much more efficiently than the PI-SDK
- Your archive server is the bread-and-butter of the operation
  - By using AF you put a layer of separation between users/apps and your most prized possession – your data
- With AF – you can standardize and customize the content for the application – and vet the data

# PI Web API – A Recommendation

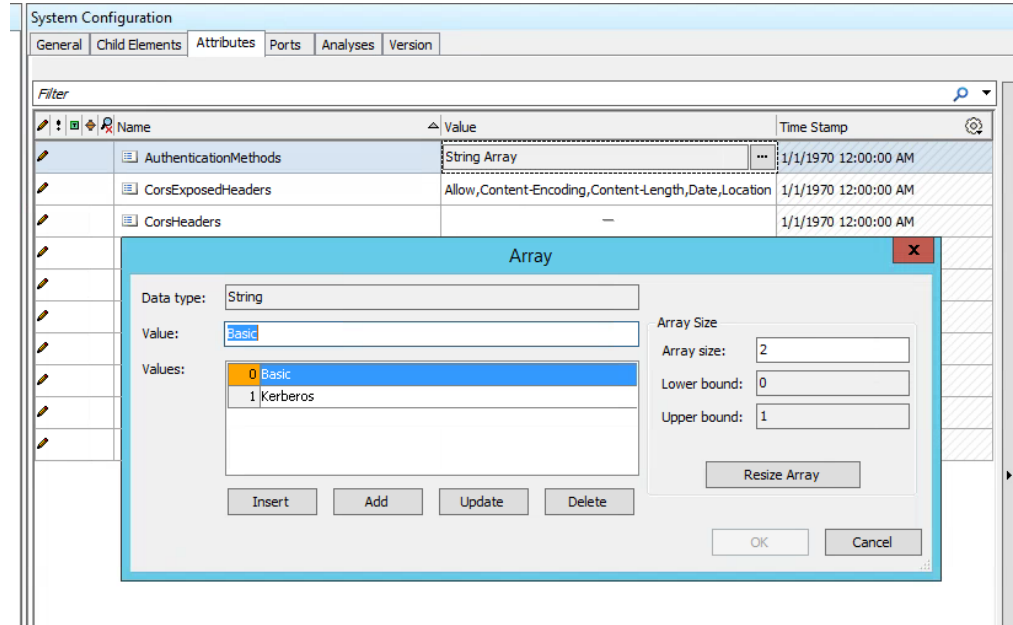
- And if your applications already involve or can use web service technologies – PI Web API is an excellent choice
  - It only works against PI AF
    - Utilizes the great features just mentioned
  - Allows read and write capabilities
  - Configurable authentication based on app host
    - Can do most secure – Kerberos – for Windows-based
    - Also can do Basic combined with https: for non-Windows

## PI Web API – A Recommendation

- Install the PI Web API concurrently on a Windows application node
- For non-Windows applications – need to install on a Windows node it can reach
  - Consider – if you are using PI Coresight, it comes with PI Web API – make it a dual use box
    - We have an actual use case in production use

# PI Web API – A Recommendation

- The configuration for the PI Web API is in your AF database called Configuration
  - Used for PI Notifications, PI Coresight and PI Analysis as well
- You can configure more than one authentication method for PI Web API
- You just need to make sure your application is able to make a choice if given multiple choices
  - Can confirm IBM MQ is one that can do this
  - PI Coresight can handle as well



# PI Web API and PI AF – Performance Proven

- More from our use case
  - PI Coresight running on Windows Server 2012 VM
  - AF model with 1,400+ attributes mapped to PI tags on PI Archive server in older domain joined by domain trust
  - Application did test of 15-minute aggregations of tag data in web query for a 1 hour interval – 4 results per attribute
  - Query returned results.... in 5 seconds
    - Application owner – sold on switch to PI Web API and PI AF!

Discussion/questions...



- [bruce.mccamant@triencon.com](mailto:bruce.mccamant@triencon.com)
- Triencon Services, Inc.
- [www.triencon.com](http://www.triencon.com)