PI System Administration: Beyond the Basics

Alex Duhig
V2021r2 / 9-Aug-2021

AVEVA Group plc
High Cross, Madingley Road
Cambridge CB3 0HB, UK
Tel +44 (0)1223 556655
Fax +44 (0)1223 556666

aveva.com

Table of contents

1 Introduction

Before reading this section, please refer to the following course YouTube video:
https://youtu.be/OE1YKqwD5i8

This course will teach PI System Administrators to upgrade, expand, and improve the reliability on the core components of a PI System, including:

- PI Interface migration and upgrade
- Cross-Domain PI Interface authentication
- PI Server Redundancy

During the class we'll cover best-practise configuration only, but will link out to documentation covering older practises and methods when necessary - although we will discuss migrating to modern best practises from outdated practises in detail.

Throughout this class we will assume you have already completed the *PI System Basics* and *PI System Administration: Basics* classes. We will not cover any background or content that has already been covered in these classes.

The course consists of text-based lessons, and video lectures. You are encouraged to follow along with all video lectures using your Learning Cloud Environment. It is intended that you follow the video lectures in-order, as a lot of configuration depends on prior lessons. Text-based lessons contain background theory, exercises, documentation and other useful information.

Lessons are categorized using the following key words:

- Hands-On: These lessons contain hands-on walkthroughs on configuring your virtual environment. Skipping these videos is not recommended, you will find that you cannot complete some later lessons without first completing the earlier ones. Throughout these sections we will link out to reference documentation. You may be examined on the content inside these references, but as the exam is open-book, you're welcome to visit these references during the exam.
- Knowledge check: Exercises to complete using your virtual environment, or quizzes to keep your skills sharp. Lessons will assume you have completed prior knowledge checks. Some knowledge checks ask you to change the configuration of your virtual environment and skipping them may result in you being unable to complete later exercises.
- Background: Readings or videos containing useful background information to assist your understanding. These lessons are optional, but encouraged as they will give you a more well-rounded understanding of the system.

In the next lesson we'll cover the Learning Cloud Environment used in class, and what we'll accomplish in this class in detail.

1.1 Background: PI System Software Used in this Class

Before reading this section, please refer to the following course YouTube video:
https://youtu.be/io_5WDrjtdY

Software	Version
PI Server	2018 SP3 Patch 3
PI Interface for OPC DA (Read Only)	2.7.1.41
PI API	2018 Patch 2
PI Interface Configuration Utility	1.5.1.10

1.2 Background: Credentials Used in this Class

Below is a table containing all credentials needed to complete the exercises in this class. The following chapters will refer to this table.

Account	Password	Reason for use
PIINT01\Student01	DescriptionExamTeacher	Logging in to the interface machine
PISCHOOL\Student01	<found in the email you received when you booted the virtual environment>	Authenticating on the domain
PIINT01\Local-PIInterface	DescriptionExamTeacher	Service account for running the opc-readonly1 and pibufss services
PISCHOOL\PIService-PIINT01	NeverUseTrusts	Service account on the domain side, for interface and buffer authentication on the domain

2 Configuring Security

2.1 Background: Upgrading and Migrating Interfaces

References (Login required):

- *What are the steps for upgrading PI Interface for OPC DA?*
- *How to move an interface to a new machine*

One of the most stressful tasks of a PI System Administrator is moving an interface from one machine to another, or upgrading one in place to a new version. Interfaces are commonly the most important link in the chain of data flow

through the system, any downtime for them typically means data loss - the worst possible outcome for a PI System Administrator.

As with all Interface administration tasks, moving or upgrading interfaces gets much less risky when you have a failover pair. If you have a pair of interfaces in failover, the update process is fairly simple:

1. Download the latest version of the interface installer from the OSIsoft Customer Portal
2. Stop the interface on one of the interface machines
3. Run the new version of the installer on the machine, and go through the installation wizard
4. Start the updated interface
5. Stop the interface on the second machine and ensure the updated interface takes over and is sending data to the Data Archive
6. Run the new version of the installer on the second machine, and go through the installation wizard
7. Start the second interface
8. Restart the first interface and ensure the second takes over and is sending data to the Data Archive

See the references above for moving an interface to new hardware, or for details on how to reduce data loss if you do not have a failover pair set up. There is one more interface maintenance task that is a little more involved. Moving from one interface type to another on the same machine. The most common reason you would do this is to move from the Read-Write version of an interface to the Read-Only version, in effort to harden security in the system. This course aims to prepare you to harden the security of your PI System - and it's exactly what we're going to do in our first hands-on.

2.2 Hands-On: Interfacing Read-only

During this section, please refer to the following course YouTube video:
https://youtu.be/g8gC-FdCX5k

In this hands-on, we will migrate a read/write version of the OPC interface to a read-only version. In doing this, we'll also get experience in interface management and maintenance. During the hands-on we use login credentials, these can be found in the previous "Credentials Used in this Class" chapter.

References:

- Documentation for the PI Interface for OPC DA can be found *here*, definitions for the parameters found in the configuration file can be found *here*.

2.2.1 Installing the Read-only Version of the Interface

1. Connect to PIINT01. NOTE: you will need to log in using the local account when connecting. If your virtual environment fails to connect, do the following:
    1. Click **Configure**
    2. Enter the following details:
        - **Domain**: PIINT01
        - **Password**: <enter the password for **PIINT01\student01** from the table in the previous "Credentials Used in this Class" chapter>

3. Click **Connect**
2. Open the shortcut to **PI Install Kits** on the desktop
3. Right click on **OPCInt_ReadOnly_2.X.X.X_.exe → Run as Administrator**
4. Click **OK** and **Next** through the installer until it completes

2.2.2 Configuring the Read-only Interface

1. On PIINT01, open Windows File Explorer and navigate to **D:\Program Files (x86)\PIPC\Interfaces\OPCInt**
2. Right click on **opcint1.bat → copy**
3. Navigate to **D:\Program Files (x86)\PIPC\Interfaces\OPCInt_ReadOnly**
4. Right click → **paste** into the folder
5. Right click on the new **opcint1.bat** file in the OPCInt_ReadOnly folder → **properties**
6. Uncheck **Read-only**
7. Click **OK**
8. right click on **opcint1.bat → edit**
9. Edit the final line of the file, replacing "**OPCInt\opcint.exe**" with "**OPCInt_ReadOnly\opcint_readonly.exe**"
10. **Save** the file and close the window
11. Open the **PI Interface Configuration Utility (ICU)** via the Start menu
12. From the top menu: **Interface → New Windows Interface Instance from BAT File...**
13. Navigate to and select **D:\Program Files (x86)\PIPC\Interfaces\OPCInt_ReadOnly\opcint1.bat**
14. Click **Open**
15. Click **OK**
16. Select the **Service** tab on the left side
17. Click **Create**
18. With the drop down menu, select the **opcint1** instance
19. Select **Yes** to the save changes prompt
20. Stop the interface with the stop button
21. With the drop down menu, select the **opcint_readonly1** instance
22. **Start** the interface with the start button
23. With the drop down menu, select the **opcint1** instance
24. From the top menu: **Interface → Delete Interface Instance...**
25. Click **No** to the prompt

2.2.3 Finishing Best Practise Configuration

1. On PIINT01, open the **ICU**, and from the drop down menu select the **opcint_readonly1 instance**
2. Select the **Disconnected Startup** tab from the left
3. Check **Enable disconnected startup (with point caching)**
4. Next to **Cache Path:** click **Browse**
5. Navigate to **D:\Program Files (x86)\PIPC\Interfaces\OPCInt_ReadOnly**
6. Click **OK**
7. Select the **Health Points** tab from the left
8. For each of the following, right click → **create**:
    - **Heartbeat**
    - **Device Status**
    - **IO Rate**
    - **Scan Class Scans Skipped.sc0**
9. Click **Apply**
10. Restart the interface with the restart button

2.3 Background: Authentication vs. Authorization

Let’s review what we know so far. In the context of the PI System:

- Authentication is the process that verifies the identity of a user or process, before allowing it to connect to the Data Archive
- Authorization is the process that determines what an application can do once connected to the Data Archive or the Asset Framework (e.g. create a PI Point, create an asset, run a backup, etc.)

We like to make an analogy of the Data Archive (or the Asset Framework) as a building. The process of authentication is like the security guard at the entrance of the building. They decide whether you should be let in. If they do let you in, you are given an access card. This access card is your authorization. It will give you access to specific rooms within the building.
In the Basics class we covered authorization, but skipped over most of authentication, mostly because only one of the types is considered best practise - so we ignored the others. In this class we're covering upgrading and moving from legacy security, so we should know an little bit more about what's out there.

There are three different methods of authentication on the Data Archive:

2.3.1 PI Mappings

PI Mappings use *Windows Integrated Security* to authenticate users on the Data Archive. With this method, users and services connect directly to the Data Archive using their Windows account. A PI Mapping grants a Windows user or group specific rights on the Data Archive by assigning a PI Identity.
This method of authentication has several advantages:

- It is the most secure
- It enables transport security (encryption in transit) of communications with the Data Archive
- It represents the least amount of maintenance for PI System administrators
- It allows users to connect directly with their Windows accounts

The recommended strategy for using PI Mappings is to create a Windows Group for each level of authentication needed on the Data Archive (e.g. one group for Read-Only users, one group for PI System Administrators, etc.), then assign a unique PI Identity to each one of these groups.
PI Mappings are created from System Management Tools, from Security > Mappings & Trusts > Mappings Tab, by pressing the New button . This will open the Add New Mapping Window

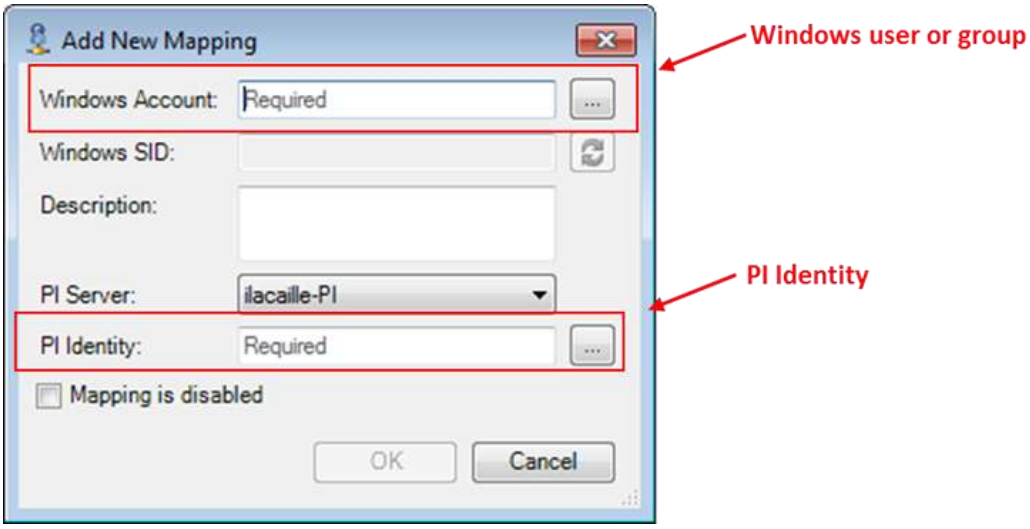

The following conditions must be true in order to use PI Mappings:

- The application must connect with **PI AFSDK (any version), PI SDK version 1.3.6 or later or the PI API for Windows Integrated Security (version 2.0.1.35 and later, released in 2016)**
- The connecting application is running on a Windows operating system

2.3.2 PI Trusts

PI Trusts should NOT be used unless it is not possible to authenticate using Windows Integrated Security. We'll discuss scenarios in which you would use trusts in the next lesson.

Note: Prior to 2016 release of the PI API for Windows Integrated Security, any applications using the PI API, such as PI Interfaces, could not use PI Mappings. Now, almost all PI Interface nodes can be upgraded to the new security model, regardless domain or workgroup configuration.

The PI Trust authentication method works by comparing some connection metadata of the connecting application to the metadata saved in PI Trusts. If the the metadata matches one of the configured trusts, the connection is allowed. No login is required by the application.
PI Trusts are created from System Management Tools, from Security > Mappings & Trusts > Trusts Tab, by pressing the arrow next to the New… button and selecting the advanced option:

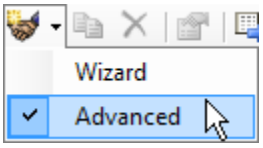

This will open the Add New Trust Window.

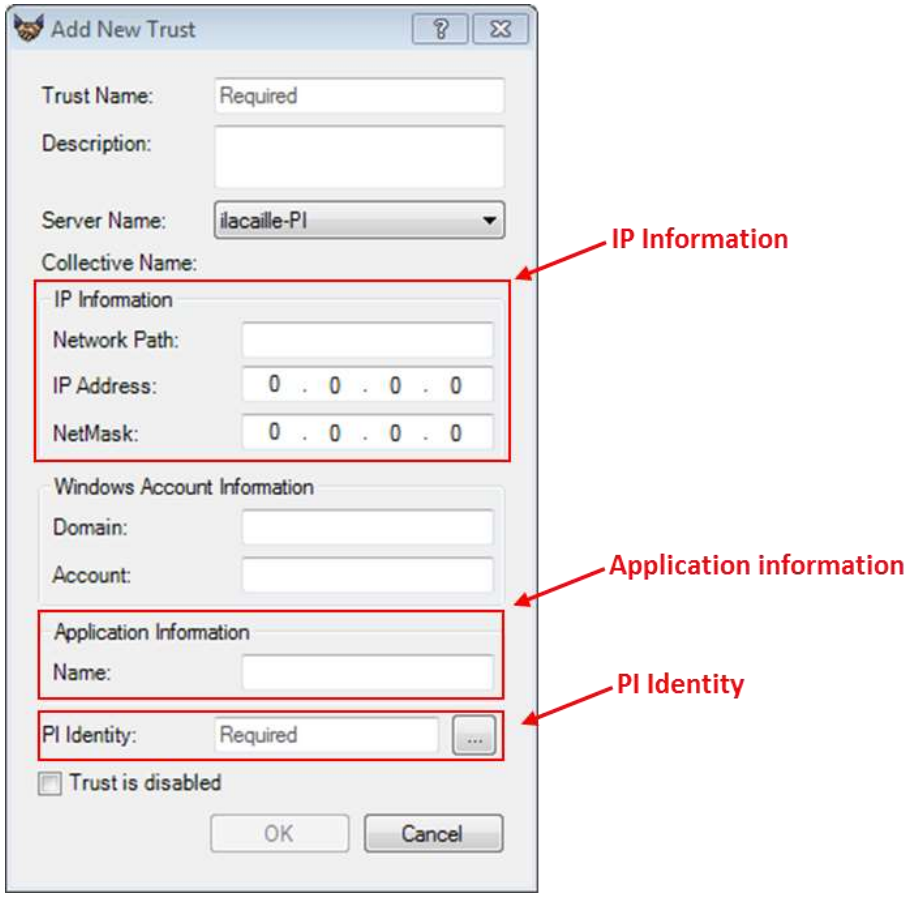

It is not necessary to fill in all of the information in this Window. OSIsoft recommends that you fill out PI Trusts using the 2+ Trust convention. This means you need to enter the following:

- **The IP Information:**
  - The Network Path (Host name or fully qualified domain name of the computer)

    OR
  - The IP Address and a NetMask of 255.255.255.255.
- **The Application Information**: The application name. Applications that connect through the PI API send an identifier called an application process name, or procname. This is a four-character string with an E appended. For example, the procname for the PI Perfmon interface is PIPeE.

2.3.3 Explicit Login

The final authentication method, Explicit Login, is not recommended in any scenario. It only exists for backwards compatibility purposes. Using this method, users login to the Data Archive directly using a PI User and a password. This method not only suffer from a lack of encryption like trusts, but also puts a burden on the administrator and users to manage separate accounts and passwords for the PI Data Archive.

2.4 Knowledge Check: PI Data Archive Authentication

In this exercise we'll take a look at authentication methods, and how to identify them in the Data Archive message logs. We'll use PISQL01 as a client machine. For now, ignore the fact that it's our SQL Server, we're going to pretend it's a

simple client PC. We're using PI SDK Utility - a simple little application installed on any PI System client PC that allows a connection test, and the viewing of log files.

Perform the following steps:

1. Connect to PISQL01,
2. Through the **Start** menu, open **PI System → PISDKUtility (64-bit)**
3. Navigate to the **PI SDK → Connections**
4. Open from the top menu: **Connections → Options**

The connection options box can be opened from many PI System clients, and shows the options the client machine uses when it connects to Data Archives. Here's where we could change the default server clients use on this machine when making their initial connection to a server, and change the preferred order of authentication protocols. Note that all settings here are for this client PC only - they're not server side settings. If you would like to push these connection settings to many client PCs at once, an article covering this can be found *here* (login required).

Note the "Protocol Order" on the right:

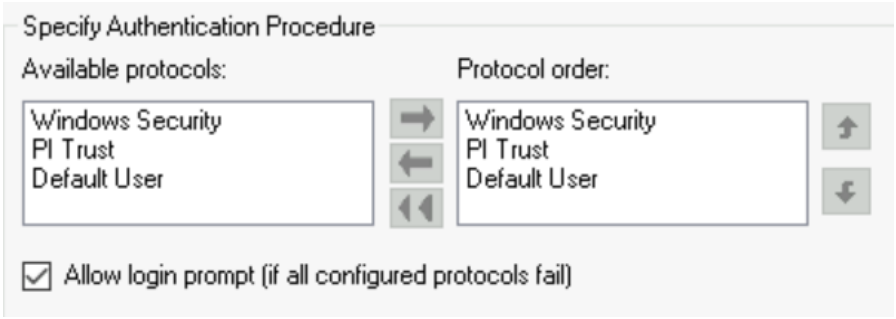

This is the default order of connections for a client. It will first try to connect with Windows Security, then with a PI Trust, then with "Default User". If all of these fail, then it pops up a login prompt. By now we know what Windows Security and PI Trusts are. Default User and Login prompt refer to the Explicit Login method described in the previous lesson. They're not used in modern systems, and should not be used due to large management overhead and security reasons. This method is turned off on the Data Archive by default.

5. Click **OK**
6. Click the Checkbox next to **PISRV01**. It should successfully connect. Note the connected user is listed as **PISCHOOL\student01 as piadmins**.
7. Open a virtual machine connection to PISRV01
8. on PISRV01, through the **Start** menu, open **PI System → PI System Management Tools**
9. Navigate to **Operation → Message Logs**
10. Click the **magnifying glas**s to retrieve log messages
11. Find the message associated with the connection from PISQL01. It should look something like this:

```
Successful login  ID: 23. Address: 192.168.0.23. Name: PISDKUtility.exe(5984):remote.
Identity List: piadmins | PIWorld. Environment Username : PISCHOOL\student01. Method:
Windows Login (SSPI,Kerberos,HMAC-SHA1-96,Kerberos AES256-CTS-HMAC-SHA1-96,256)
```

This message indicates that a mapping was used to authenticate this connection, and the user connected was PISCHOOL\student01. Let's try using a trust.

12. on PISRV01, inside **SMT**, navigate to **Security → Mappings and Trusts**
13. Select the **Trusts** tab
14. Click the drop-down next to the **New...** button on the top left, and select **Advanced**

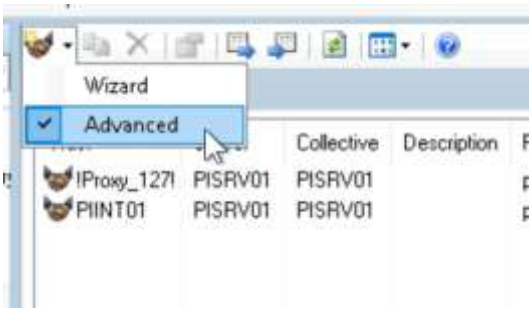

15. Enter the following:
    - **Trust Name:** PISDKUtility - PISQL01
    - **Network Path:** PISQL01
    - **Application Information → Name:** PISDKUtility.exe
    - **PI Identity:** piadmins
16. Click **OK**

We just created a 2+ trust. A trust that uses at least two pieces of connection metadata to authenticate incoming connections. Let's test it out.

17. On PISQL01, in **PI SDK Utility**, uncheck the box next to **PISRV01**, then check it again
18. Note that we still connect as a Windows User
19. Check the log on PISRV01 inside **SMT**, note that we're still using a mapping rather than a trust

What's wrong? Has our configuration failed? It hasn't, it's just that the connection protocol order prefers Windows authentication, so will connect using that if possible. Let's change it.

20. On PISQL01, in **PI SDK Utility**, select **Connections → Options**
21. Select **PI Trust** inside **Protocol Order**, then click the **Up** arrow. This should promote it to be first on the list
22. Click **OK**, then uncheck and recheck the connection box
23. Note that the Connected User is now **PI Trust as piadmins**
24. Check the log in **SMT** on PISRV01 (ensure you click the magnifying class "retrieve logs" button again), note that we're connecting using a trust:

```
Successful login ID: 16. Address: 192.168.0.23. Host: . Name:
PISDKUtility.exe(5984):remote. User: piadmins. OSUser: . Trust: PISDKUtility - PISQL01
```

One important thing to note: the user here is what will be printed on subsequent log messages when this user performs tasks. This is one of the primary reasons you should **never** use trusts for user connections. There's no way to audit or identify who made changes to teh system if multiple users are trusted to the same identity.

Now... let's remove this trust. I'm uncomfortable just looking at it.

25. On PISRV01, in SMT, navigate to **Security → Mappings & Trusts**, **Trusts** tab

26. Right click on **PISDKUtility - PISQL01 → Delete**
27. Click **Yes**
28. On PISQL01, in **PI SDK Utility**, uncheck the box next to **PISRV01**, then check it again. Note that we connect again with our Windows user
29. Retrieve the logs on PISRV01 again. We see two messages now:

```
Unsuccessful login  ID: 20. Address: 192.168.0.23. Name: PISDKUtility.exe(5984):remote.
Credentials used: PISCHOOL\student01. Method: Trust. Error: [-10413] No trust relation
for this request

Successful login  ID: 20. Address: 192.168.0.23. Name: PISDKUtility.exe(5984):remote.
Identity List: piadmins | PIWorld. Environment Username : PISCHOOL\student01. Method:
Windows Login (SSPI,Kerberos,HMAC-SHA1-96,Kerberos AES256-CTS-HMAC-SHA1-96,256)
```

Our PISDKUtility attempted to connect using a trust, which failed. Then it tried using Windows authentication, which succeeded. This is because we had trust first in our protocol order.

30. On PISQL01, in **PI SDK Utility**, go to **Connections → Options** and change the protocol order back so **Windows Security** is at the top

Most clients connecting to the Data Archive use this protocol order to decide how they connect, and it can be changed through this **Connections → Options** window. The notable exception for this is interfaces. Interfaces ignore this setting, and will behave like so:

- If PI API for Windows Integrated Security is installed on the interface machine, the interface will **only** connect using Windows Security
- If PI API for Windows Integrated Security is **not** installed on the interface machine, the interface will **only** connect using Trusts

In the next lesson we'll talk about possible scenarios that may force you to use trusts over Windows Security.

2.5 Background: So do I Ever Use Trusts?

You may be asking yourself "Do I ever use trusts then?"

Not unless you can't help it.

There are some specific situations where you **must** use trusts:

- If you're using a PI Interface installed on a UNIX operating system
- If you're using a very old PI API application that was written and published before the year 2016, and have lost the source code so cannot update it to run on more recent versions. Note that we're specifically saying "PI API" here. PI SDK has been a much more common library to develop applications with since the year 2009.

That's it. Only two real dealbreakers. They come up very rarely nowadays. If you must use trusts, you should follow the practises set out in *this article* (login required) under the section "More Specific Trusts – Multiple Credentials". These are sometimes called "2+ Trusts".

So why are so many organisations still using trusts? The following is a summary of reasons why a lot of organisations stick to trusts and will not move to mappings:

Inexperience or lack of initiative to move to more modern standards

This is the most straightforward to fix - attend a class like this, and ensure that either you're never in a situation where you need to move (you use best practises during installation), or move your system over after completion of this course.

Organisation policies may not allow accounts with non-expiring passwords to be created

A lot of organisations have a blanket rule, stating that accounts with passwords must expire after a time period. This makes it infeasible to manage the different service account passwords on individual interface machines. It's hard to work around this other than trying to change the policy itself. Policies like this are outdated and go against modern password practises like those suggested by *NIST*. If your organisation has policies like this, start a conversation with your IT department asking why they are in place, and if they should be reviewed with modern guidelines in mind.

All in all, the risk caused by an old password being used by an interface is far less risky than using a trust. If the two are compared, there's no contest as to which is preferable when looking objectively.

Administration accounts on interface machines are difficult to manage

Imagine you have your interface on a process control network with its own domain, pushing data to a Data Archive on the corporate network. This works just fine, and is a very common topology. The configuration we went through in this course is what is recommended; use Windows Credential Manager to have the interface log in to the Data Archive using a domain account.

The complication comes when you think about your administration team. Your admin team would have individual accounts on the process control network (PCN) domain, and different individual accounts on the corporate domain. When an administrator needs to make a change to an interface, they must start up the ICU, which needs to log in to the PI Data Archive with close-to admin rights. Your administrators would need to run the CMDKEY command, mapping their PCN account to their corporate account after they log in.

Now imagine you have hundreds of interfaces, and a large group of administrators. It becomes frustrating and time consuming for administrators to map credentials whenever they log in to any interface node. You may also feel uncomfortable having admins map their corporate credentials from a different domain.

There's no easy solution for this problem. Security is always a balancing act between security and convenience. If the above manual mapping won't work for you, other options along the security/convenience spectrum are:

- Use trusts, following the practises set out in *this article* (login required) under the section "More Specific Trusts – Multiple Credentials". These are sometimes called "2+ Trusts". Going with this option sacrifices the the benefits of over-the-wire encryption for interface connections, and increases attack surface as trusts cannot be disabled.
- Instead of using the ICU, manually manage interface and buffer configuration files. Administrators would need to regularly reference the documentation for the interface they are configuring, and *documentation* and *articles* on buffer configuration. Going with this option
- Create a shared account on the PCN domain, limited to only being able to run the ICU. Create a CMDKEY entry for this account to a corporate domain account with minimum privileges that the ICU needs.

2.6 Hands-On: Configuring Cross-Domain Security

During this section, please refer to the following course YouTube video:
https://youtu.be/6cfAX1RCkL4

In this hands-on, we will configure secure cross-domain access for an interface on a machine that is not a member of the same domain as the PI Data Archive. During the hands-on we use several different login credentials, these can be found in a table on the online version of this lesson. When this exercise mentions “the table above”, it is referring to the table on the online version of the lesson.

2.6.1 Configuring Authentication for Administrators

1. Connect to PIINT01. NOTE: you will need to log in using the local account when connecting. If your virtual environment fails to connect, do the following:
    1. Click **Configure**
    2. Enter the Domain: **PIINT01**
    3. Enter the Password as the password that appears in the previous "Credentials Used in this Class" chapter for the **local account** PIINT01\Student01
    4. Click **Connect**
2. Open a Windows Command Prompt
3. Type the following command. **Note**: replace the text <YOUR PASSWORD> with the password for the **domain account** PISCHOOL\Student01. You will find the password for this account in the previous "Credentials Used in this Class" chapter.

```
CMDKEY /add:PISRV01.PISCHOOL.INT /user:PISCHOOL\student01 /pass:<YOUR PASSWORD>
```

4. Press the **Enter** key to execute the command

2.6.2 Preparing the Data Archive

1. Connect to PISRV01
2. Open **PI System Management Tools (SMT)** from the Start menu
3. Navigate to **Security → Identities, Users & Groups**
4. While on the **PI Identities** tab, click the **New Identity** button on the top left of the pane
5. **Name** this identity "PIINT01-OPC"
6. Click **Create**
7. Open **Microsoft Excel** from the Start menu
8. If a license warning appears, close it. This warning appears because the virtual environment cloning process makes Excel think the environment has new hardware. We should be fine riding out the grace period here.
9. Click **Blank workbook**
10. Click the **PI Bulider** tab
11. **PI Points → Find PI Points...**
12. Type "pointsource:PIINT01-OPC" into the **search** box
13. Click **Search**
14. Click **OK**
15. Click **Clear All**

16. Select the checkmark next to the **Security** heading
17. Click **OK**
18. Add the following characters to the end of each cell under the **datasecurity** header: " | PIINT01-OPC: A(w)"
19. Add the following characters to the end of each cell under the **ptsecurity** header: " | PIINT01-OPC: A(r)"
20. Click **Publish**
21. Select **Edit Only**
22. Click **OK**
23. Click **Close**
24. Open **SMT** and navigate to **Security → Mappings & Trusts**
25. Click **Add New Mapping**
26. Write in **Windows Account**: "PISCHOOL\PIService-PIINT01"
27. Write in **PI Identity**: "PIINT01-OPC"

2.6.3 Configuring Authentication for the Interface

1. Connect to PIINT01
2. Open a **command prompt** from the start menu
3. Right click on the command prompt icon on the taskbar
4. Hold the **shift** key, then right click on the text "Command Prompt" in the pop up menu
5. Click **Run as different user**
6. For the credentials, use:
    - **User**: PIINT01\Local-PIInterface
7. For the **password**, enter the corresponding password for this account in the previous "Credentials Used in this Class" chapter
8. Click **OK**
9. Enter the following command, ensuring you replace <OBTAIN FROM ABOVE TABLE> with the password for the PISCHOOL\PIService-PIINT01 account in the previous "Credentials Used in this Class" chapter:

```
CMDKEY /add:PISRV01.PISCHOOL.INT /user:PISCHOOL\PIService-PIINT01 /pass:<OBTAIN FROM
ABOVE TABLE>
```

10. Press the **Enter** key to execute the command
11. Click **Start**
12. type "lusrmgr.msc" and press the Enter key
13. Navigate to **Groups → PI Buffering Administrators**
14. Add the "local-PIInterface" user account to the group
15. Navigate to **Windows Services** by clicking the services button on the taskbar
16. Right click on **PI Buffer Subsystem → Properties**
17. on the **Log On** tab, do the following:
    - Select the radio button next to **This account:**
    - Enter as the account **name**: ".\Local-PIInterface"
    - Enter the corresponding **password** for PIINT01\Local-PIInterface in the previous "Credentials Used in this Class" chapter
18. Click **OK**
19. Click **OK**
20. Repeat steps 15 through 17 for the **PI-opcint_readonly1** service
21. Right click on **PI Buffer Subsystem** service **→ Restart**
22. Select **Yes**

2.6.4 Installing PI API for WIS

1. Connect to PIINT01
2. Open the **PI Install Kits** shortcut on the desktop
3. Right click on **PIAPI.X.X.X.exe → Run as Administrator**
4. Click **OK**
5. Click **Install**
6. Click **OK** or **Next** on all further prompts during the install
7. If asked to reboot, say **Yes**

2.6.5 Finally saying goodbye to trusts

1. Connect to PISRV01
2. Open **SMT**
3. Navigate to **Security → Security Settings**
4. Drag the slider to the top
5. Click **Save**

3 High Availability and Redundancy

In this chater we'll discuss high availability and redundancy for Data Archives, Asset Framework, and how these redundancy features interact with UNIINT Interface Failover. We'll go on to discuss

3.1 Background: How PI Server High Availability Works

Before reading this section, please refer to the following course YouTube video:
https://youtu.be/Kolx8Q1ohBI

References:

- You can read more about Data Archive and AF high availability *here*.
- For directions on installing the backend Asset Framework (AF) database on a SQL Server availability group (highly available SQL Servers) see the documentation *here*.
- In this course we only cover the recommended AF High Availability (HA) solution - Load Balanced AF Servers. For full documentation including other options for AF HA see the documentation *here*.

3.2 Hands-On: Installing AF with a Remote SQL Backend

During this section, please refer to the following course YouTube video:
https://youtu.be/6cfAX1RCkL4

In this hands-on we'll install Asset Framework and set up a backend database on a remote server.

References:

- SQL Server versions supported by AF are detailed *here*.
- Installation of AF on SQL backends with different highly available architectures can be found here - *Availability groups*, *Mirrored*, *Failover Clusters*. Directions on installation can be found in the chapters of these documentation sections.

3.2.1 Asset Framework Installation

1. Connect to PISRV01
2. Open the **PI Install Kits** shortcut on the desktop
3. right click on **PI-Server_XXXX....exe → Run as Administrator**
4. Select **Modify**, then click **Next**
5. Check the box next to **AF Server**, then click **Next**
6. In the the **SQL Server Connection** box, type "PISQL01\sqlexpress"
7. Uncheck **AF SQL script Execution**
8. Uncheck **Validate connection to SQL Server...**
9. Click **Next**
10. Click **Select**, then click **OK**
11. Click **Next**
12. Type the following under the **Account Name** Columns
    - **PI AF Application Service**: "PISCHOOL\SVC-PIAF$"
    - **PISQL DAS (RTQP Engine)**: "PISCHOOL\SVC-PIRTQP$"
13. Click **Next**, then click **Modify**
14. Wait for the install to finish, then click **Close**

3.2.2 Copying the SQL Scripts

1. Connect to PISRV01
2. Open **Windows File Explorer** and navigate to **D:\Program Files\PIPC\AF**
3. Right click on the **SQL** folder → **Copy**
4. Open the **Shared Folder** shortcut on the desktop
5. Right click and **paste** inside the **Shared Folder** to copy the **SQL** folder inside
6. Connect to PISQL01
7. Open the **Shared Folder** shortcut on the desktop
8. Right click on the **SQL** folder → **Copy**
9. Right click somewhere on the desktop → **Paste** to copy the **SQL** folder to the desktop

3.2.3 Running the SQL Scripts

1. Connect to PISQL01
2. Run a **Windows Command Prompt**
3. execute the following commands, pressing the **Enter** key between each:

```
cd Desktop
cd SQL
go.bat PISQL01\SQLEXPRESS PIFD
```

4. Wait for the script to finish running

3.2.4 Allowing AF Server Access

1. Connect to PISQL01
2. Click the **Start** menu, then open **Microsoft SQL Server Tools 18 → Microsoft SQL Server Management Studio**
3. Click **Connect**
4. Expand **Security**
5. right click on **Logins → New Login...**
6. Click **Search...**
7. Click Locations...
8. Select the **Entire Directory** then click **OK**
9. Click **Object Types...**
10. Check the box next to **Service Accounts** then click **OK**
11. Enter the object name: "PISCHOOL\SVC-PIAF$", then click **OK**
12. Click **User Mapping**
13. Check the box next to **PIFD**
14. Check the boxes next to:
    - **db_AFQueryEngine**
    - **db_AFServer**
      [Note - this step is actually incorrect. We'll fix this in the next exercise. Assign these permissions now anyway so you're ready to fix them the next exercise!]
15. Click **OK**

3.3 Knowledge Check: Correcting Permissions for RTQP

Oh no! We made a mistake when we were configuring security for the PIFD database. Completely intentional, I swear.

Our RTQP service runs on the PISCHOOL\SVC-PIRTQP$ account, but we configured security for the PISCHOOL\SVC-PIAF$ account. If anyone tried to use our RTQP service right now, they'd fail to connect. If you're interested in exactly what this service does, feel free to read through the documentation *here*, or take a look at our course on *Exposing PI Data with the PI SQL Framework*. In simple terms, it exposes the AF server as if it were a SQL server, allowing users to run SQL queries to extract PI Data.

The RTQP service requires direct access to the AF backend database, with slightly different permissions than the AF Server. When users run queries against the engine, it executes queries against the SQL backend database. It will also connect to the Data Archive and retrieve any data it needs to service user queries. You should never execute queries against the AF SQL backend database itself to retrieve your AF hierarchy or production data, this should all be done through the RTQP or similar services. The only reasons to connect to the backend database directly is during this initial setup, or to back up or move the database.

3.3.1 Exercise - Fix Your Mistake

Your task is simple. Remove the db_AFQueryEngine access for the SVC-PIAF$ account and give the SVC-PIRTQP$ account this access instead.

[Try to complete this on your own before following the solution]

The solution can be found below.

3.3.2 Solution - Fix Your Mistake

1. Connect to PISQL01
2. Click the **Start** menu, then open **Microsoft SQL Server Tools 18 → Microsoft SQL Server Management Studio**
3. Click **Connect**
4. Expand **Security**
5. Expand **Logins**
6. Double click on **PISCHOOL\SVC-PIAF$**
7. Select **User Mapping**
8. Click on the text **PIFD**
9. Uncheck **db_AFQueryEngine**
10. Click **OK**
11. Right click on **Logins → New Login...**
12. Click **Search...**
13. Click **Locations...**
14. Select the **Entire Directory** then click **OK**
15. Click **Object Types...**
16. Check the box next to **Service Accounts** then click **OK**
17. Enter the object name: "PISCHOOL\SVC-PIRTQP$", then click **OK**
18. Click **User Mapping**
19. Check the box next to **PIFD**
20. Check the box next to **db_AFQueryEngine**
21. Click **OK**

3.4 Background: Preparing for a Highly Available PI Server

We're about to install a highly available PI Server - meaning installing and configuring both the Data Archive and Asset Framework components on PISRV02, and setting them up to be a highly available pair with their counterparts on PISRV01. In the real world, you'll need to do some preparation before you install a secondary PI Server:

1. Obtain a new license file for both your Primary Data Archive and Secondary Data Archive(s)

If you're expanding your architecture from a single Data Archive to a collective and adding to your license, you'll need to update the license on your primary with the new license that supports the second collective member. You can generate a new Data Archive license on the *customer portal* by uploading a Machine Signature File from your primary archive and downloading a new license file. The package downloaded will contain a pilicense.dat (which should be used on your primary Data Archive machine) and temporarylicense.dat (which should be used on any secondaries). More information can be found in *the documentation*.

2. Update your existing Data Archive license on your primary archive machine, or install a PI Server using this new license
3. Ensure appropriate firewall ports are open on the primary and secondary machines. A full list of ports can be found *here*.
4. Ensure your user account has the permissions needed to set up the collective. You can read the requirements *here*.

3.5 Hands-On: Implementing PI Server High Availability

During this section, please refer to the following course YouTube video:
https://youtu.be/_ykO3UMvFyw

In this hands-on we'll configure redundancy for both the Data Archive and AF Server

References:

- In the video we mentioned pushing out known servers table updates to all users on your network. For more information this, see *this article* (login required).
- For the privilege required for creating a Data Archive collective, see the documentation *here*.
- For more information about managing collectives with PI Collective Manager, see the documentation *here*.
- We mentioned creating analytics, events and notifications during the video. These topics are covered in courses in our *Power User* training path, and official documentation for these features can be found *here*.

3.5.1 Installing the Second PI Server

1. Connect to PISRV02
2. Open the **PI Install Kits** shortcut on the desktop
3. right click on **PI-Server_XXXX....exe → Run as Administrator**
4. Click **Next**
5. Check the box next to **all** four components
6. Change the **PI Data Archive directory** to "D:\Program Files\PI"
7. Click **Next**
8. Change the **SQL Server Connection to** "PISQL01\sqlexpress"
9. Uncheck both checkboxes:
    - **AF SQL database scripts**
    - **AF SQL script execution**
10. Click **Next**
11. Change the following installation options:
    - **License directory**: "D:\PI Install Kits\Training License"
    - **Historical Archives:** "D:\PI\arc\"
    - **Future Archives:** "D:\PI\arc\future"
    - **Event Queues:** "D:\PI\queue"
12. Click **Next**
13. Enter the following:
    - **SMTP Server*:** "PISQL01.PISCHOOL.INT"
    - **From Email*:** "notifications@pischool.int"
14. Click **Next**
15. Click **Select** then **OK**
16. Click **Next**
17. Type in the following service account names:
    - **PI AF Application Service:** "PISCHOOL\SVC-PIAF$"
    - **PI Notifications Service:** "PISCHOOL\SVC-PINOTIF$"
    - **PI Analysis Service:** "PISCHOOL\SVC-PIANALYT$"
    - **PI SQL DAS (RTQP Engine):** "PISCHOOL\SVC-PIRTQP$"

18. Click **Next** then click **Install**
19. When the installer finishes, click **Close**

3.5.2 Configuring AF Redundancy

1. Connect to PISQL01 (note, this could be done from any machine. We're running these actions on PISQL01 just to display we can do it remotely)
2. Open **PI System Explorer** from the Start menu
3. Click **Yes**
4. Enter a **Host: "**PIAF"
5. Click **Connect**, then **OK**
6. Click **Rename**
7. Write **New Name:** PIAF
8. Click **OK**, then **OK**

3.5.3 Configuring Data Archive Redundancy

1. Connect to PISRV01
2. From the **Start** menu, open **PI System → PI Collective Manager**
3. From the top menu, select **File → Connections...**
4. Right click the white space under **PISRV01 → Add Server...**
5. Type **Network Node**: "PISRV02"
6. With your browser cloud environment connection, connect to PISRV02
7. On PISRV02, Open **SMT** from the **Start** menu
8. Navigate to **Security → Mappings & Trusts**
9. Type the following:
    - **Windows Account:** "PISCHOOL\student01"
    - **PI Identity:** "piadmins"
10. Click **Create**
11. Switch your cloud environment connection back to PISRV01
12. On PISRV01, click **OK** on the **Add Server** window again
13. Click **Close**
14. From the top menu **File → Create New Collective...**
15. Check both checkboxes on the page, then click **Next**
16. Select **An existing server that contains data**, then click **Next**
17. For Collective Primary, select **PISRV01**, then click **Next**
18. Under Secondary Servers, select Server: **PISRV01**, then click **Add**
19. Click **Next**
20. Click **Next** on the **Select Archives** page
21. Click **Next** on the **Select Backup Location** page
22. Click **Next** on the **Verify Selections** page
23. Once the process completes, click **Finish**

3.6 Knowledge Check: Authenticating our Interface

After our installation, we found that our interface buffer couldn't write to our secondary PI Data Archive, because it didn't log in to PISRV02 with the correct credentials.

3.6.1 Exercise - Fix the account credentials

Again, our task is simple. Make PIINT01's pibufss and interface services use the correct credentials when logging in to the secondary Data Archive.

The solution can be found below.

[Try to complete this on your own before following the solution]

3.6.2 Solution - Fix the account credentials

1. Connect to PIINT01
2. Open a **command prompt** from the **start** menu
3. Right click on the **command prompt icon** on the taskbar
4. Hold the shift key, then right click on the text **Command Prompt** in the pop up menu
5. Click **Run as different user**
6. For the credentials, use:
    - **User**: PIINT01\Local-PIInterface
    - For the **password**, enter the corresponding password for this account in the previous "Credentials Used in this Class" chapter
7. Click **OK**
8. Enter the following command, ensuring you replace <OBTAIN FROM TABLE> with the password for the PISCHOOL\PIService-PIINT01 account in the table found in the previous "Credentials Used in this Class" chapter:

```
CMDKEY /add:PISRV02.PISCHOOL.INT /user:PISCHOOL\PIService-PIINT01 /pass:<OBTAIN FROM
TABLE>
```

9. Press the **Enter** key to execute the command
10. Navigate to **Windows Services** by clicking the services button on the taskbar
11. Right click on **PI Buffer Subsystem** service → **Restart**
12. Select **Yes**
13. Confirm data is flowing to both Data Archive servers using the **Current Values** plugin of **PI SMT** on **PISRV01**

3.7 Knowledge Check: Configure Automatic Data Archive Backups

Read through *PI Data Archive collectives and backups*. This will give you a good overview of backup strategy with Data Archive Collectives, including best practises.

Note: We haven't mentioned backups for a redundant set of AF servers, but the procedure is the same as with a single AF Server - simply back up the PIFD backend database.

3.7.1 Exercise - Configure Backups

Let's assume our PISRV01 and PISRV02 are geographically separated and we'd like to keep independent backups of both our primary and secondary Your task is to configure the scheduled automatic backup task using the documentation above to ensure you're following best practises. We'd like the backups to be stored in E:\pibackup on each machine.

[Try to complete this on your own before following the solution]

The solution can be found below.

3.7.2 Solution - Configure Backups

1. Connect to PISRV01
2. Open a **Windows Command Prompt**
3. Run the commands:

```
cd /d "%piserver%adm"
pibackup e:\pibackup -install
```

4. Connect to PISRV02
5. Open a **Windows Command Prompt**
6. Run the commands:

```
cd /d "%piserver%adm"
pibackup e:\pibackup -install
```

7. Open the **start** menu and type "Task Scheduler", and run **Windows Task Scheduler**
8. Navigate to **Task Scheduler Library**
9. Right click on **PI Server Backup → Properties**
10. On the **Triggers** tab, select the **Daily schedule** and click **Edit...**
11. Change the time to 3:45AM, then click **OK**
12. Click **OK**

3.8 Knowledge Check: Service Permissions and Buffering

We have our highly available system up and running - but we don't have all our services writing and reading from our servers yet. You may remember in the basics class we had to configure access to the Data Archive and AF Server for our Analysis and RTQP services. We're going to do that again here, but showcase how security and buffering setup interact with a Data Archive collective.

3.8.1 Exercise - Let the services in!

Our PISRV02 machine runs the **PI Analysis Service** and **PI SQL Data Access Server (RTQP Engine)**, and **PI Notifications Service** services, running under their respective GMSAs. Create mappings on your Data Archive collective, giving them the permission they need to access all data on the servers. You will need to configure the security for these services in line with the documentation starting at these sections: *Analysis , RTQP* and *Notifications*.

Note: Step 1b in the Analysis security configuration procedure has you configure read/write access to the PIPoint database, so the Analysis service can create its own points. While this is functionality the Analysis service has, it is best security practise to only give the service read access to the database, and have your administrators create points for the service using PI System Explorer or PI Builder. Allowing the Analysis service to create its own points effectively allows anyone that can edit the AF database the ability to create points, which can enable accidental (or intentional) changes to PI Points through the service.

You do not need to configure security for the AF side of things - when we ran the PI Server installer and specified which accounts we'd like to run these services, the installer automatically configured these mappings.

HINT: you only need to modify identities, mappings, and tag configuration on the primary, all changes will be automatically synchronised over to the secondary.

[Try to complete this on your own before following the solution]

The solution can be found below.

3.8.2 Solution - Let the services in!

1. Connect to PISRV01
2. Open **PI System Management Tools**
3. Check the box next to **PISRV01** in the top left corner
4. Go to **Security → Identities Users and Groups**
5. Create three new identities:
    - SVC-PIRTQP
    - SVC-PIANALYT
    - SVC-PINOTIF
6. Go to **Security → Mappings and Trusts**
7. Create three new mappings:
    - Mapping PISCHOOL\SVC-PIRTQP$ to SVC-PIRTQP
    - Mapping PISCHOOL\SVC-PIANALYT$ to SVC-PIANALYT
    - Mapping PISCHOOL\SVC-PINOTIF$ to SVC-PINOTIF
8. Go to **Security → Database Security**
9. Double click on **PIPOINT**
10. Add your three new identities to the list and ensure they have Read access

Optional: If we were to strictly follow best practises, we would now open Excel and use PI Builder to give the new identities Read access to PointSecurity and DataSecurity for all points. Modifying the PIPOINT database security doesn't modify the security for all points, but only the default security for new points. However, we haven't disabled PI World access to our server, so the services already have this access to all points. We'll skip these steps, as the Basics class covered them at length.

3.8.3 Exercise - Configure Buffering to a Data Archive collective

PISRV02 runs the Analysis service, and now we have a collective - if it had permission to, it would only write to tags on the primary server. If we'd configured buffering before we created the collective, it would have automatically started buffering to both servers... but we didn't. Open the Buffer Manager from PI System Explorer and set up buffering to both members of our Data Archive collective.

HINT: Refer to the lesson "The Bufferer, the Better" in PI System Administration: Basics for inspiration on all the configuration you need to do.

[Try to complete this on your own before following the solution]

The solution can be found below.

3.8.4 Solution - Configure Buffering to a Data Archive collective

1. Connect to **PISRV01**
2. Select **Tools > Buffering Manager**
3. Click **Yes** to the question asking if you'd like to configure buffering
4. Click **Continue with Configuration**
5. Change the Buffer location to **D:\Buffer**
6. Click **Next**
7. Exit the Wizard

Now we've configured buffering, we need to make sure it connects to the PI Data Archive using an account that can write to analysis points. In an earlier exercise, we configured the SVC-PIANALYT group managed service account to have these privileges.

8. Open the **Windows Services** panel by clicking on the icon for it in the taskbar
9. Modify the **PI Buffer Subsystem** service so it logs in as the **PISCHOOL\SVC-PIANALYT** account (hint: this account is a group-managed service account so you may need to do special things with dollar signs and passwords. If you don't know what to do, the configuration is very similar to the configuration we did in the interface buffering video, but with a different account this time).
10. Create a service dependency for **PI Analysis Service** on **PI Buffer Subsystem** to ensure that PI Buffer Subsystem starts first on machine start. To do this, run **Windows command prompt** with the icon in the taskbar and execute the following command:

```
SC CONFIG PIAnalysisManager depend= pibufss
```

11. Open **Computer Management** from the Windows Start menu. Navigate to **System Tools → Local Users and Groups → Groups**
12. Double click on the **PI Buffering Administrators** group, and add the **SVC-PIANALYT** account to the group
13. Use the **Windows Services** panel to restart the **PI Buffer Subsystem service**. It should warn you that it will restart the PI Analysis service as well.
14. Close and reopen **Buffering Manager**. The status indicator should be green if you have configured buffering correctly.

3.9 Background: When do we Add More Servers?

During this section, please refer to the following course YouTube video:
https://youtu.be/CpS62V-CCzM

References:

- Documentation on available connectors can be found under the "Connectors" section of *this page*
- For directions on how to move the Analysis service to a new machine, see *here*
- Documentation on installing the Analysis service on a Windows Failover Cluster can be found *here*
- Documentation on installing PI Vision can be found *here*

- The Developer Learning Path can be found *here*.
- Documentation on moving PI Server components in depth can be found under the *Installation and Upgrade section* of the documentation. Each component has a section on upgrading and moving with procedures covering standalone and redundant configurations.

4 Final Exam

The final exam in this course is taken online. Please check the course listing online for more details.