

CA BEST PRACTICES

CA IT Process Automation Manager Best Practices

Guidelines for Continuous and
High Availability

DRAFT DOCUMENT – [FEEDBACK](#) WELCOME!

LEGAL NOTICE

This publication is based on current information and resource allocations as of its date of publication and is subject to change or withdrawal by CA at any time without notice. The information in this publication could include typographical errors or technical inaccuracies. CA may make modifications to any CA product, software program, method or procedure described in this publication at any time without notice.

Any reference in this publication to non-CA products and non-CA websites are provided for convenience only and shall not serve as CA's endorsement of such products or websites. Your use of such products, websites, and any information regarding such products or any materials provided with such products or at such websites shall be at your own risk.

Notwithstanding anything in this publication to the contrary, this publication shall not (i) constitute product documentation or specifications under any existing or future written license agreement or services agreement relating to any CA software product, or be subject to any warranty set forth in any such written agreement; (ii) serve to affect the rights and/or obligations of CA or its licensees under any existing or future written license agreement or services agreement relating to any CA software product; or (iii) serve to amend any product documentation or specifications for any CA software product. The development, release and timing of any features or functionality described in this publication remain at CA's sole discretion.

The information in this publication is based upon CA's experiences with the referenced software products in a variety of development and customer environments. Past performance of the software products in such development and customer environments is not indicative of the future performance of such software products in identical, similar or different environments. CA does not warrant that the software products will operate as specifically set forth in this publication. CA will support only the referenced products in accordance with (i) the documentation and specifications provided with the referenced product, and (ii) CA's then-current maintenance and support policy for the referenced product.

Certain information in this publication may outline CA's general product direction. All information in this publication is for your informational purposes only and may not be incorporated into any contract. CA assumes no responsibility for the accuracy or completeness of the information. To the extent permitted by applicable law, CA provides this document "AS IS" without warranty of any kind, including, without limitation, any implied warranties of merchantability, fitness for a particular purpose, or non-infringement. In no event will CA be liable for any loss or damage, direct or indirect, from the use of this document, including, without limitation, lost profits, lost investment, business interruption, goodwill or lost data, even if CA is expressly advised of the possibility of such damages.

COPYRIGHT LICENSE AND NOTICE:

This publication may contain sample application programming code and/or language which illustrate programming techniques on various operating systems. Notwithstanding anything to the contrary contained in this publication, such sample code does not constitute licensed products or software under any CA license or services agreement. You may copy, modify and use this sample code for the purposes of performing the installation methods and routines described in this document. These samples have not been tested. CA does not make, and you may not rely on, any promise, express or implied, of reliability, serviceability or function of the sample code.

Copyright © 2009 CA. All rights reserved. All trademarks, trade names, service marks and logos referenced herein belong to their respective companies. Microsoft product screen shots reprinted with permission from Microsoft Corporation.

TITLE AND PUBLICATION DATE:

CA IT Process Automation Manager Best Practices – Guidelines for Continuous and High Availability
Publication Date: April 2010

Last Update: April 12, 2010

ACKNOWLEDGEMENTS

Principal Authors and Technical Editors

Farid Charkhian
Anders Magnusson
Yatin Dawada
Terry Pisauro

The principal authors and CA would like to thank the following contributors:

Anil Parthasarathy

CA PRODUCT REFERENCES

This document references the following CA products:

- CA IT Process Automation Manager™ (CA IT PAM)
- CA Embedded Entitlements Manager (CA EEM)
- CA Enterprise Log Manager

FEEDBACK

This is a draft document – feedback is welcome! Please email us at impcdfedback@ca.com to share your feedback on this publication. Please include the title of this publication in the subject of your email response. For technical assistance with a CA product, please contact CA Technical Support at <http://ca.com/support>. For assistance with support specific to Japanese operating systems, please contact CA at <http://www.casupport.jp>.

Contents

Chapter 1: Introduction	7
Chapter 2: Defining Failover for Continuous Availability	9
Application Failure.....	9
CA IT PAM Component Failure	11
Agent Failover	12
Orchestrator Failover.....	13
Database Failure.....	15
Hardware Failure	17
Security Component Failure	18
Subnet/Network failure.....	18
Site Failure	18
Disaster Recover Example.....	21
Chapter 3: Using Clustering for High Availability	23
Before You Begin	23
Requirements for Domain Orchestrator Nodes	23
Servers for Shared Components	24
Install/Configure Supporting Components	25
Install the Apache HTTP Server.....	25
Install CA EEM	27
Install CA IT PAM Domain Orchestrator on Primary Node	29
Install CA IT PAM Prerequisites	29
Install CA IT PAM Domain Orchestrator on Primary Node	30
Launch the CA IT PAM Console	33
Install Domain Orchestrator on Secondary Node	34
Verify Cluster Functionality.....	36
Perform a Basic Test of the Loadbalancer and Cluster Nodes	36
Test Failover of CA IT PAM while a Process is Being Edited	37
Test Failover of CA IT PAM with an Active Workflow Process	39
Appendix A: Advanced Configurations of Apache HTTP Server and CA IT PAM	47
Mod_jk.so Module	47
Contents of the mod-jk.conf file.....	47
Contents of the workers.properties file	48
Customizing the workers.properties File.....	49
Contents of uriworkermap.properties file for a Domain Orchestrator.....	50
Contents of uriworkermap.properties for non-Domain Orchestrators.....	51



Chapter 1: Introduction

CA IT Process Automation Manager (CA IT PAM) provides a centralized and structured approach to operations management by enabling you to define, build, orchestrate, manage, and report on automated processes spanning across different teams and roles in your organization. By automating routine administrative tasks, CA IT PAM improves operational efficiency and incident response handling, and ensures best practice and regulatory controls compliance.

This document is one in a series of papers providing best practices for making the most of your CA IT PAM implementation. The focus of this paper is on planning for both High Availability and Continuous Availability – also known as “Disaster Recovery”. It provides guidelines for planning your implementation so that key functions remain available and accessible in the event the hardware providing those functions becomes unavailable. In addition to supporting availability of CA IT PAM, the use of clustering solutions also improves the scalability of CA IT PAM.

Additional best practices documents are also available from the CA IT PAM Best Practices pages which are accessible from the following link:

https://support.ca.com/phpdocs/0/common/impcd/r11/Catalyst/ITPAM_Frame_sc.htm

Chapter 2: Defining Failover for Continuous Availability

Planning for Disaster Recovery – or Failover – can help ensure continuity of business in the face of the unexpected – from the loss of a supporting application to the loss of a server and even the loss of an entire site. CA IT PAM supports failover for the following situations:

- Application failure
- CA IT PAM component failure – Orchestrators (both Domain and non-Domain), Agent, Proxy Agent
- Database failure
- Security component failure
- Hardware failure
- Subnet/network failure
- Site failure

Application Failure

Application failure refers to failures in processes run by the CA IT PAM workflow engine. CA IT PAM is designed to handle these failures and will do the following:

- Automatically retry non destructive operations
- Suspend all other operations
- Execute exception handlers when specified within the process flow
- Recover the current state of operations that completed during the failure state

If the failure path cannot be determined, CA IT PAM will log errors and wait until the error is corrected or remediated. Suspension of current operations requires Operator intervention. New process invocations will continue as defined.

Application failover planning occurs at the process level. The process designer can take appropriate steps to handle this including: retries, sending alerts (writing to loggers, email, SNMP traps, etc.) to other applications (like Event Managers, etc.) or to individual people (Admins, Operators, Level 1 Support personnel, etc.). This can be built by using various process constructs provided within CA IT PAM (JMX, SNMP, SMTP connectors, exception handlers, swim lanes, etc.).

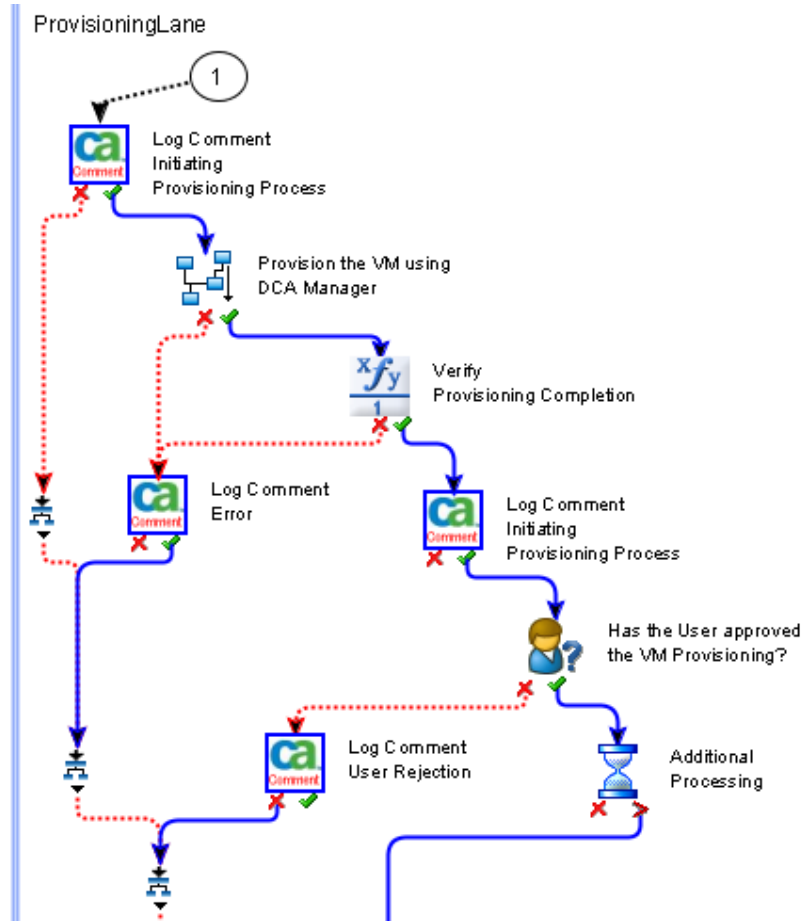
If the process is active and the Agent goes down, the process status will not change until the agent is restarted. If the process is mission critical then, as part of process design, you can send Events to the management system to check for completion. If the completion code is not received within the expected time frame the necessary alerts can then be generated.



Redundancy on the Touchpoint/Agent level can be implemented by assigning multiple agents to a touch point. When this is done the operator will normally route all traffic to the associated agent with the lowest priority and only use agents with higher priority when other agents aren't available.

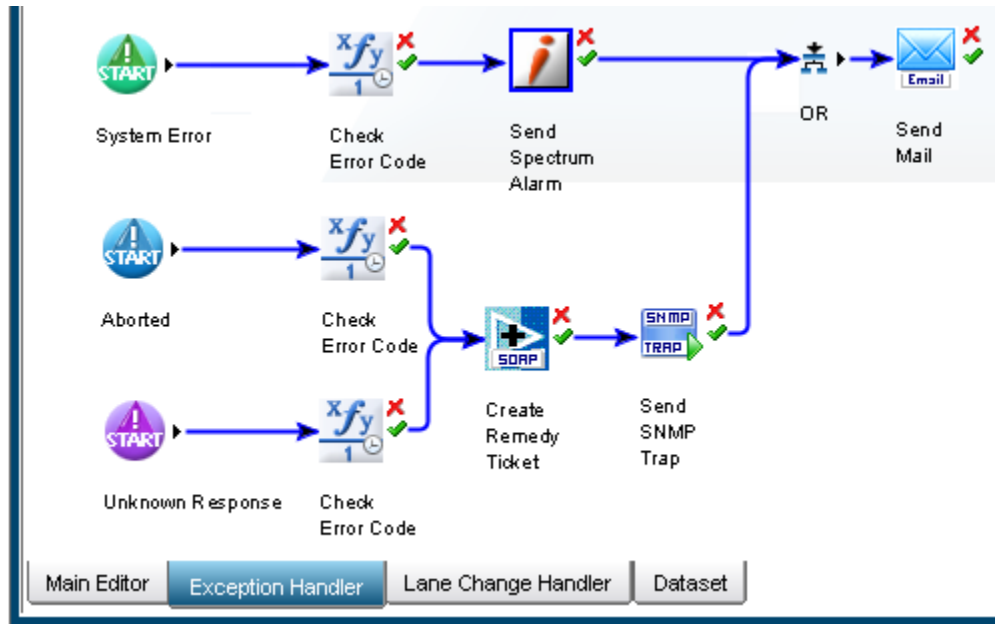
While the Agent is down any new processes will automatically go to an Agent with lower priority or, if the running process is resumed, it will also run on the Active agent with the lower priority.

Here you can see an example of application failure planning:



This example highlights the importance of correctly logging any "milestones" in the Process as well as any unexpected results, errors, rejected procedures and so on. To visually separate paths that are used to track exceptions it is a good idea to change the link properties (color and style) so that they are easy to identify.

It is also important to correctly define the exception handler so that all otherwise uncaught exceptions are handled correctly. An example of this can be seen below:



As you can see in this example the exception handler can trap “System Errors”, “Aborted Processes” and “Unknown Responses” and, depending on the type of exception, take different paths, analyze the error code and notify the operation using all methods that are available within CA IT PAM.

Common reasons for the various exception types are:

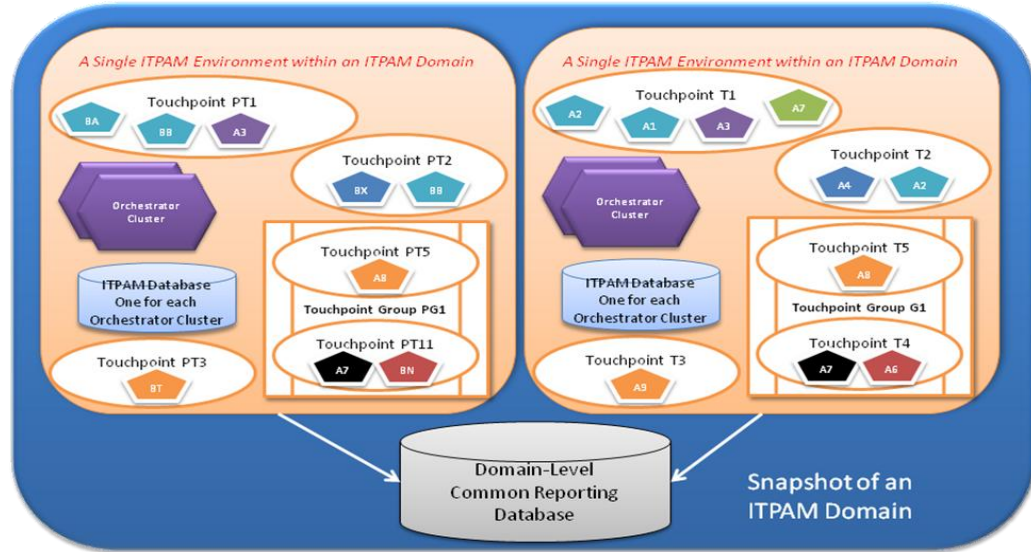
- System Exception
Usually occurs with any type of communication failure. For example when the process contains an incorrect Touchpoint name or refers to agent that can’t be accessed.
- Unidentified Response
Occurs when no output connector responds to the operator’s response.
- Aborted
Triggered when a user or administrator abort the process.

CA IT PAM Component Failure

CA IT PAM can be configured to manage the failure of the following infrastructure components automatically:

- Orchestrators (both Domain and non-Domain)
- Agents
- Proxy Agents

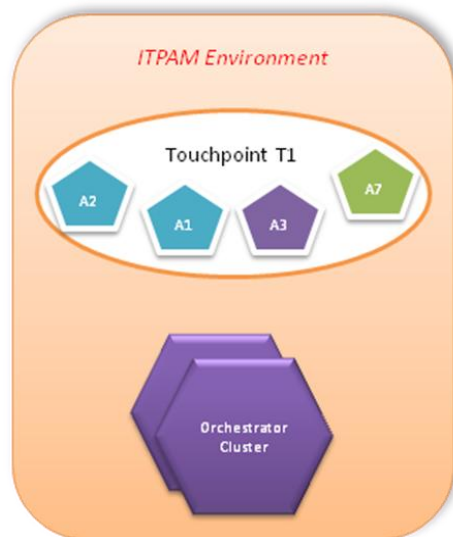
Here you can see an example of an IT PAM configuration:



Agent Failover

Agent failover, as well as load balancing, is managed through the concept of Touchpoints. A Touchpoint is a naming abstraction for one or more Agents and a single agent may be associated with multiple Environments through Touchpoints. One or more Agents running on one or more pieces of physical hardware may belong to a single Touchpoint. The Domain Orchestrator manages the state of the Agents. Individual Orchestrators inherit the state of the Agents within their Environment from the Domain Orchestrator. If an Agent is active, all Touchpoints associated with that Agent will be Available. Agents within a Touchpoint may be prioritized for Load-Balancing or Failover and Touchpoints, themselves, can be grouped for functionality.

Consider the following example:



In this example, Touchpoint T1 is configured to have the following four CA IT PAM Agents:

- A2, running on machine A2
- A1, running on machine A1
- A3, running on machine A3
- A7, running on machine A7

The priorities are as follows:

- A1 and A2 share the same priority
- A3 has a different priority
- A7 has a different priority

When an operation is invoked on Touchpoint T1, the **OR** condition is honored. In other words, only one of the Agents in Touchpoint T1 will execute the operation. Given this:

- A1 and A2 share the same priority and hence load-balance
- A3 and A7 are configured for failover

When multiple concurrent operations are initiated on T1, all of the work is shared between Agents A1 and A2 and no work is sent to Agents A3 and A7. Only when both A1 and A2 are down, will the operations be sent to the next lower priority Agent and so forth down the line until at least one Agent in the Touchpoint executes the operation. When **no** Agents are available within a Touchpoint, the operation is failed by the Orchestrator. Note that multiple Orchestrators (clustered or standalone) will share the Touchpoints within an IT PAM environment.

CA IT PAM Agent failures are managed by Touchpoint Mappings to Agents. For most critical Agents, generally those that handle host-based processing, it is possible to configure a Touchpoint with the Agent with higher priority and a Proxy Agent on the machine (using SSH) as the failover Agent with a lower priority. For Agents that invoke operations that handle remote connectivity, like Web Services, multiple Agents may be associated with a Touchpoint to handle failover.

Failure for CA IT PAM Proxy Agents are managed in a fashion similar to CA IT PAM Agents

Orchestrator Failover

Failover support for CA IT PAM Orchestrators is provided through clustering which also provides load-balancing for backend components. Consider the following example:



Orchestrators use JBoss (J2EE container) to implement clustering. Orchestrators provide two interfaces to services

- Web Interface
- Service Interface

The Service Interface is managed by clustering which is transparent to users. Failover management is completed automatically and does not require any additional administration, configuration and management.

The Web Interface requires an external solution. Orchestrators use Web Service Interfaces for communications. Clients use JNLP to connect to the Orchestrators within an environment. In a cluster each Orchestrator has its own Web Service Interface. Using an external load balancer ensures both failover and load-balancing for the CA IT PAM Orchestrator Cluster. CA IT PAM has been tested using Apache Web Server for external Load-Balancing to handle Failover for Web Interface. The load balancer will:

- Maintain its own Cluster Membership Information
- Validate its Membership List
- Redirect incoming Web Service Connections appropriately to the Cluster Members
- Manage redirection of priorities, etc.
- Additionally the following features are supported with AJP (Apache Jserv Protocol):
 - > Support Sessions (*Sticky Sessions*)
 - > Cache Connections
 - > Cache Security Credentials

An Orchestrator cluster may have one or more Orchestrators on different machines in different locations. The first Orchestrator that is active will automatically assume the role of the “master” within the cluster. The remaining Orchestrators (if any) will register as “slaves” within the cluster.

Orchestrators do not support priorities. Within a Cluster, all registered Orchestrators will automatically load-balance. The master Orchestrator will handle the following services, in addition to normal services shared among all Orchestrators within the cluster:

- Data Source Management Services
- Naming Services (*Application Service Lookups*)
- Queuing Services

When a master Orchestrator fails, any of the slave Orchestrators within the cluster will automatically assume the role of the master and set up these services. During the transition from old master to new master the following will occur:

- New incoming requests to the old master will respond with “Server not Available”
- Internal requests to the old master will be retried until they are successfully handled
- All currently running operations that are not local to the old master will be successfully recovered

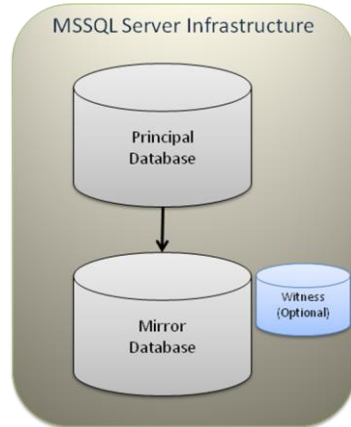
Operations local to the old master will be recovered when the old master comes back online.

Database Failure

Options for database failure planning will depend on the best practices established at your site. Most companies typically use the database replication solution provided by the database vendor (such as Oracle, MS, MySQL, etc.) or other hardware data replication solutions provided by their storage vendor (EMC, etc.). CA IT PAM will use a logical database connection, such that, if the Primary database fails over, CA IT PAM automatically switches to the Secondary database (as configured for your site’s best practice).

The Common Reporting Database may also be configured for Disaster Recovery. For example, you can use a vendor provided solution to set up mirroring from Primary to Secondary Database.

Consider the following example, using MS SQL Server:



In this example three machines are required for automatic failover:

- Principal
- Mirror
- Witness

Two Machines/Servers are required for participating in manual failover:

- Principal
- Mirror

This configuration is supported only in MS-SQL Server Editions (Standard, Developer and Enterprise) and requires configuration changes on the SQL Server end after CA IT PAM has been installed. CA IT PAM should be shut down before implementing database failover changes.

Setup Mode:

- High Availability
- FULL Transaction Safety
- Synchronous Transfer Mechanism
- Quorum Required
- Witness Server Required
- Automatic Failover

The following procedures were tested on MS SQL Server 2005 SP2.

Mirroring Setup:

- Prepare the Mirror by restoring (with **NORECOVERY** Option) the Principal backup for these databases
- The Principal and Mirror Instances of MS SQL Server require physical machines
- The Witness Instance may run on a Virtual Machine
- Ensure DNS is setup right for these three machines. Fully Qualified Domain Names are mandatory for Mirroring

- Network Firewalls, TCP/IP Filtering, etc. are turned off to ensure communications between these machines
- Set up the MS SQL Server to run as a service using either an Admin Account within the same work group or a Domain Account. *Local System Account will not work!*
- All three machines must use the same User Credentials for Mirroring
- Follow Microsoft's documentation on setting up the Mirroring Configuration for:
 - > ITPAM Library Database
 - > ITPAM Queues Database
 - > ITPAM Reporting Database
- Validate the Mirroring Setup with simple tests

The following changes should be made on the Orchestrator machine to set up IT PAM for MS SQL Server Failover

- Edit the OasisConfig.properties file to add the following keywords:

```
oasis.database.additionalparamurl  
oasis.reporting.database.additionalparamurl
```

The values for these properties depend on the Mirror Server Name FDQN. For example, if the mirror server is named "SERVER01.ABC.COM" then the value would be:

```
oasis.reporting.database.additionalparamurl=;encrypt=false;integratedSecurity=false;failoverPartner=SERVER01.ABC.com
```

- Restart the Orchestrator

The Orchestrator is now configured for automated failover of the MS SQL Server database configured with mirroring for the IT PAM databases.

Note: In addition to the MS SQL scenario described above, CA IT PAM also supports other database engines, such as Oracle and MySQL. These require other configuration steps to ensure that all CA IT PAM databases (Library, Queues and Reporting) are secured. Those details are beyond the scope of this document.

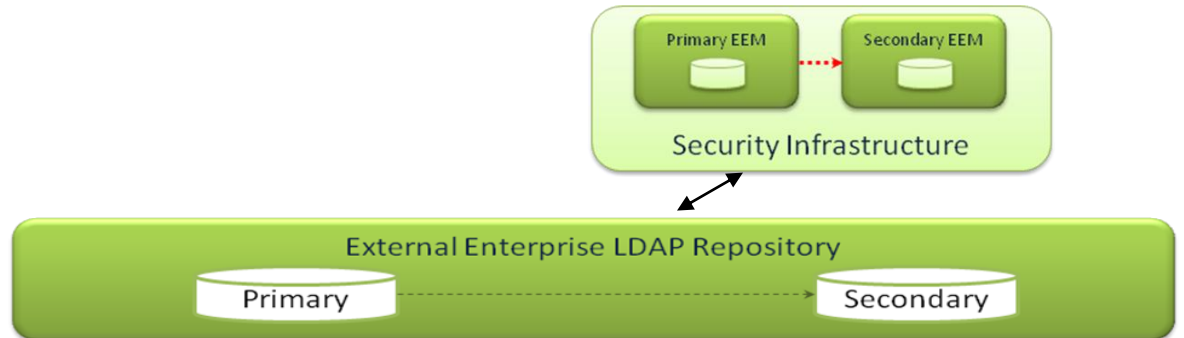
Hardware Failure

CA IT PAM supports High Availability to offset hardware failure. The particular architecture configuration will vary based on the criticality of each server group on which CA IT PAM runs. For example, financial systems typically have fully fault-tolerant automatic failover and CA IT PAM can be configured to do the same using Touchpoints.

For Orchestrators, CA IT PAM can be configured for clustering to handle failover and DR scenarios. This ensures that the Orchestrators can back each other up and can also provide auto synchronization of resources, so that in case an Orchestrator becomes unavailable, the other Orchestrators within the cluster transparently fail over without the need for any manual intervention.

Security Component Failure

Here you can see the Security infrastructure used by CA IT PAM:



EEM can be installed on each cluster node with slave (secondary domain). However, this does require replication of application data in the EEM data store. A resource kit is available to enable EEM HA when running in a Microsoft Cluster Service (MCSC) environment. This kit will automatically replicate application data on each EEM Server running on the cluster.

Each Secondary Domain (slave) will connect to its local EEM.

For more information on EEM configuration see Chapter 7 in the EEM *Getting Started Guide*. Note that both the Primary and Secondary EEM servers must be running the same version of CA EEM and must be synchronized on time. The EEM version must be 8.4 or later and the Application Name and Certificate Password must be the same for both the Primary and Secondary EEM.

Subnet/Network failure

Network failover is handled in a similar fashion to Hardware failure. From the perspective of CA IT PAM, network router failure or a hardware failure are identical. Whether the Agent is unreachable because the Server is down or because the network connectivity between the Orchestrator and the Agent is down, the Orchestrator will retry the same operation on another backup Agent as previously configured.

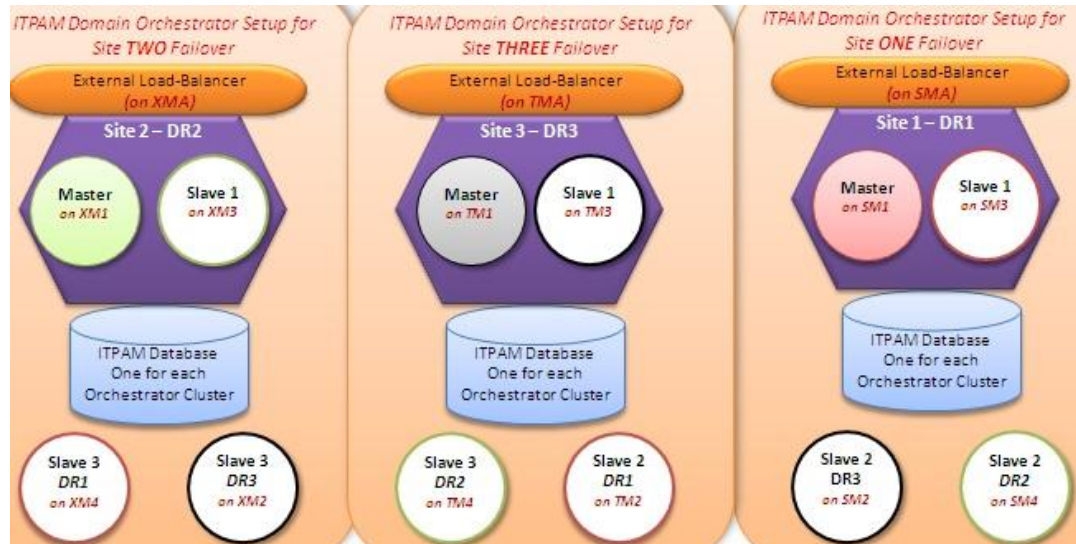
If the failure path cannot be determined, CA IT PAM will log errors and wait until the error is rectified or remediated.

Site Failure

When a Site (or Datacenter) is unreachable, you simply have to start up the Secondary Orchestrator node (provided one has been configured and appropriately synchronized with the Primary Orchestrator).

The Secondary Orchestrator node will automatically start processing upon successful activation. In this scenario, you can failover to the DR site using simple instructions and with minimal disruption to your business.

Each Site (or Datacenter) should be setup as a standalone Domain Cluster and, additionally, at least one Domain Orchestrator should be installed at the Disaster Recovery location for the Site as a “Slave Node” within the main cluster for the Site. This Slave Node will always be down after initial configuration and setup and will never participate in any of the regular load-balancing and updates from the Master Node within the Site Cluster. However, it should be periodically started to obtain its updates from the Master Node of the Site Cluster during a maintenance window and then shut down to ensure data and resource synchronization.



In the example above each site has a domain called DR1, DR2 and DR3, and each one of these has a Master and a Slave locally and Slave 2 and Slave 3 that is located in the other sites.

Each site also needs access to a highly available database solution where all CA IT PAM domain databases are available even if a site goes down. Depending on the database solution the details of how to implement this varies, however the key is that the data from each site is replicated and accessible from the alternative sites.

When the Site is down due to any reasons, the following information will have to be transferred to the Slave Node at the other Disaster Recovery location if the slave node was also down for an extended period and hence could not update automatically:

- Configuration File (Domain.xml)
- System Resources and Executables (C2O Repository)

The Slave Node will then start at the Disaster Recovery Location and automatically assume the role of the Master Node for the Site, since the entire Site Cluster is down. The Slave Node will be able to connect to various Agents and continue its processing without any downtime. Once the Site is active the following will occur:

- The Site Cluster is restarted
- The Slaves Nodes at the Site Cluster will mirror all the required data and resources
- The Slave Node is shutdown
- A new Master emerges from the Slave Nodes at the Site

The external load balancer will ensure failover and load-balancing for a CA IT PAM Orchestrator Cluster at the Site. Every Site can alternatively host a Slave Node for Disaster Recovery of other Sites

This setup can be implemented for other critical Orchestrators as well. Each Site will automatically back up as a Disaster Recovery Site for another Site. Multiple Sites may be setup as Disaster Recovery locations for a single Site and thus provide additional Disaster Recovery capacity. Since each Cluster share the same Unique Identifications and Signatures, they will be able to assume the role of the new Master with minimal disruption to business.

Periodic Maintenance Windows are required for synchronizing the DR Site Slaves to the Site Master Cluster Configuration and Resources. The Information to be synchronized must be manually copied over to the new Master Cluster if the Slave Node is not auto-synchronized with the Master Node. The manual transfer of data can be scripted (in other words, included as CA IT PAM processes which are periodically scheduled) by the local CA IT PAM Administrator to ensure DR synchronization, however, since a CA IT PAM Cluster supports heterogeneous operating systems, depending on the OS of the current Master and the DR backup, these copy instructions may require syntax changes.

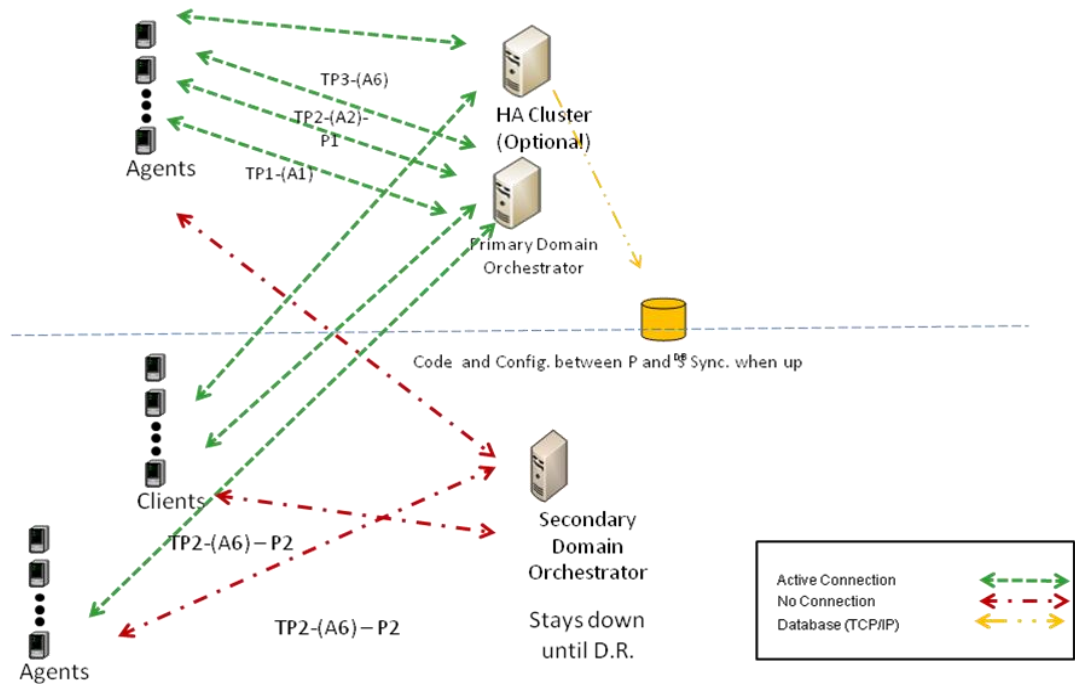
Periodic Disaster Recovery Tests should be conducted quarterly – or according to site best practice recommendations – in order to ensure that the Site Failover Procedures are well documented and adhered to by the Support Personnel.

Note: Keep in mind that cluster nodes used as failover servers for a Domain Orchestrator require full connectivity to all Agents and Orchestrators in the domain. Communication between Orchestrators and Agents requires two ports; intra-cluster communication requires several ports, and makes use of a multicast address to maintain the state of the nodes in the logical/physical cluster. Please refer to the OasisConfig.properties for the list of all ports required.

Disaster Recover Example

Here you can see an example of a typical CA IT PAM Deployment Architecture:

Normal Operations



In this example there are two main Domain Orchestrators: one residing in the "Primary Site" and the other in the "Secondary Site". The Secondary Orchestrator functions as a backup and is down until a disaster event occurs.

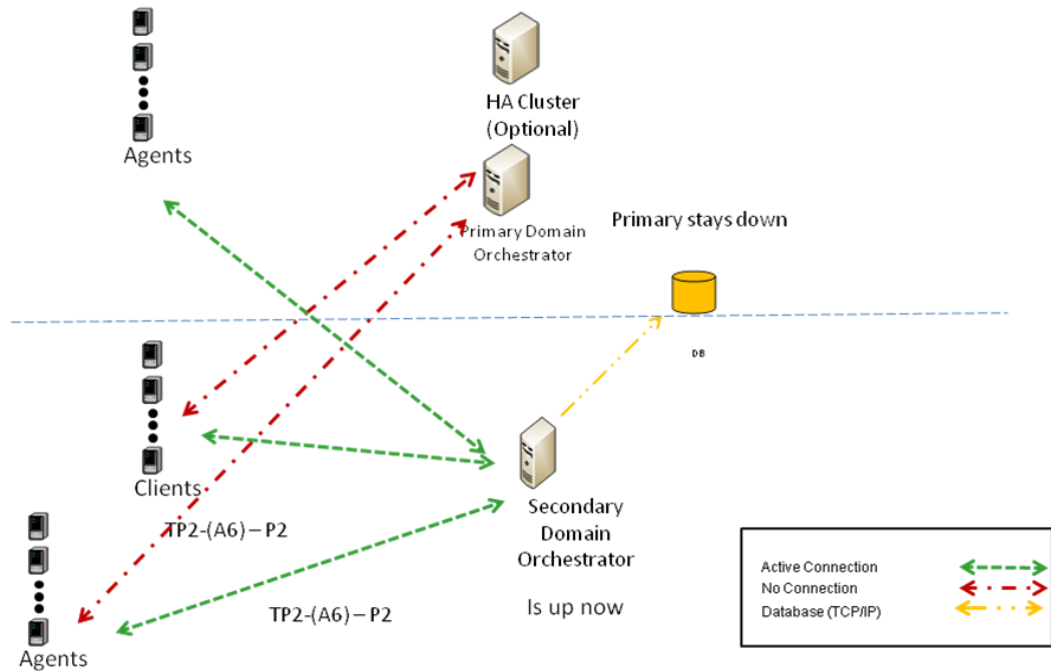
In this setup, the disaster recovery procedure for the database is independent of that of the Orchestrators, and it is conceivable that the primary Orchestrator would fail without the database failing. Both Orchestrators will point to the same 'logical' database, but that database will be replicated using available RDBMS methods that could vary from Log Shipping, to Mirroring to Clustering.

If additional capacity or High Availability of the Primary/secondary Orchestrators is needed, cluster nodes can be added to these Orchestrators as well.

The deployed CA IT PAM code and configurations, including the Domain.xml and user scripts, are automatically kept in synch between all Orchestrators, as long as those Orchestrators are up and running. The configuration and code of Orchestrators that are turned off in normal operations (such as those placed at the DR site) need to be synchronized through some other means.

In the example there are three agents, Agent 1, Agent 2 and Agent 3, at the Primary Site. Each of these agents is associated with a Touchpoint – TP1, TP2 and TP3, respectively. In addition, Agent 2 has a priority 1. In the Secondary Site, there is an Agent 6, which is also associated with TouchPoint TP2, but with Priority 2. Here you can see what would happen in a failover:

After Disaster Recovery



In the event of a failover, the Secondary Domain Orchestrator is brought up manually, by following a simple procedure. If there are Touchpoints that were shared across the two sites, that have the same agent, the pending tasks are now resumed, but is now overseen by the Secondary Domain Orchestrator.

Chapter 3: Using Clustering for High Availability

This chapter describes how the Domain Orchestrator can be easily clustered over a secondary Domain Server to enhance scalability and high availability.

Note: The following steps detail the installation of a clustered Domain Orchestrator, using CA Embedded Entitlements Manager (CA EEM) as a Security Server, Microsoft SQL Server 2005 as the datastore and Apache HTTP Server as a load balancer.

Before You Begin

If you are familiar with installing and working with CA IT PAM you will notice that installing a secondary clustered node is very simple - once you have the supporting infrastructure correctly set up. A basic standalone Domain Orchestrator is often co-located with the database server and CA EEM, however in a clustered environment these shared components need to be located on a separate server. In addition you will also need to have access to a loadbalancer. In our scenario we will use the Apache HTTP Server with the `mod_jk` module as the loadbalancer.

Following are the requirements for installing an Orchestrator on a cluster:

- An Apache JServ Protocol 1.3 (apj13) capable load balancer pre-configured in your environment. The Apache HTTP Server is a freely available load balancer. You will need "mod_jk.so" (Apache to Tomcat connector) installed on an Apache Server as your load balancer. For more information, refer to <http://apache.org>.
- Hostname / IP Address of your load balancer and port on which it is running.
- You will need to specify a unique node name (Worker Node) for the Orchestrator instance that you are going to install to maintain sticky sessions. Sticky sessions will improve the efficiency of a clustered environment.
- Depending on the load balancer you used, the Worker Node must be configured on your load balancer before the Orchestrator fully becomes an operational cluster. You can change the port where the IT PAM Orchestrator is listening for ajp communication in the `/installationDir/Orchestrator/c2o/.conf/OasisConfig.properties` file by changing "protocol.tomcat.connector.ajp.port" value after you complete your Orchestrator installation. Note that, if you make any changes to your `.properties` file, you must restart your Orchestrator.

Note: Since securing these shared components is not a focus of this document, we are going to locate all of them on one separate third node. Typically, in a production environment, these components would be implemented in highly available environment.

Based on this we will need a minimum of three servers for our environment, one for each orchestrator node and one for the shared components. This is in addition to the Server/Machines that should be managed by CA IT PAM Agents.

Requirements for Domain Orchestrator Nodes



In our example the Domain Orchestrators nodes each require the following:

- Real or Virtual systems running Microsoft Windows Server 2003 with current maintenance applied
- Java 1.6 JDK installed and integrated with the web browser
- The installation media for CA IT PAM 2.1 SP3. This should be obtained (downloadable from <http://support.ca.com>) and staged to be accessible from the primary node.

Installation of the operating system and the Java JDK are not covered in this document.

Servers for Shared Components

In our example all shared components, such as the Load balancer, Microsoft SQL Server and EEM, will be installed on one separate server. The prerequisites for this server are:

- Real or Virtual systems running Microsoft Windows Server 2003 with current maintenance applied. If IIS, or any other application binding to port 80, is installed, its service (“World Wide Web Publishing Service”) should be disabled.
- Microsoft SQL Server 2005 (any edition) with mixed mode security enabled and with current maintenance applied

Note: To ensure high availability for the complete CA IT PAM solution in a production environment it is recommended that Microsoft SQL is implemented in a highly available cluster.

- The installation media for CA IT PAM 2.1 SP3 (with CA EEM) should be accessible
- Access to the code for an Apache JServ Protocol 1.3 (apj13) capable load balancer. In our example we will use:
 - > Apache HTTP Server (apache_2.2.14-win32-x86-openssl-0.9.8k.msi or later)
 - > Apache to Tomcat Connector (mod_jk-1.2.28-httpd-2.2.3.so or later). Ensure the connector version matches the Apache HTTP Server version.

Both of these components can be downloaded from <http://apache.org>. The direct link to the HTTP Server is <http://httpd.apache.org/download.cgi> and the link to the connector is: <http://tomcat.apache.org/download-connectors.cgi>.

Installation of the operating system and the Microsoft SQL Server are not covered in this document.

Install/Configure Supporting Components

This section provides detailed instructions on how to install and configure the required shared components on a separate shared server.

Install the Apache HTTP Server

The first step is to install the loadbalancer on the shared server. To do this:

1. Ensure that you have access to the "Apache HTTP Server" (apache_2.2.14-win32-x86-openssl-0.9.8k.msi or later) and the corresponding "Apache to Tomcat Connector" (mod_jk-1.2.28-httpd-2.2.3.so). See the section "Servers for Shared Components" on page 24 for details on this.
2. If Microsoft IIS (or any other webservice binding to port 80) is installed ensure that this service is uninstalled or disabled.
3. Double-click on the install package "apache_2.2.14-win32-x86-openssl-0.9.8k.msi"
The Welcome dialog displays.
4. Click Next
The License Agreement dialog displays.
5. Review and select "I accept the terms in the license agreement" and click Next.
A Read This First dialog displays.
6. Click Next
The Server Information dialog appears.
7. Verify that the correct server information appears for Network Domain, Server Name and Administrator's Email address and that the option "for All Users, on port 80, as a Service" is selected.

Note that the Server Name should reflect the name of the server on which the shared components will be installed.
8. Click Next.
9. Select the Setup Type "Typical" and click Next.
10. You will be prompted to confirm the install destination folder.
11. Click Next.
12. Click Install on the "Ready to Install the Program" panel.
13. Click Finish on the "Installation Wizard Completed" screen.
14. Verify that the Apache HTTP server is working by opening the page <http://localhost/>
If everything is working so far the page should display "It works!"

Configure the Apache HTTP Server to work as a Load Balancer

At this point you have successfully installed the Apache HTTP Server. The next step is to configure it as a loadbalancer for CA IT PAM. To do this:

1. Rename the previously downloaded file "mod_jk-1.2.28-httpd-2.2.3.so" to "mod_jk.so" and copy it to the following folder:

```
${APACHE_INSTALLED_LOC}\modules
```

If you have not already downloaded the mod_jk.so file you can download the binary corresponding to the OS and installed version of Apache HTTP Server from <http://tomcat.apache.org/download-connectors.cgi>.

In our example we are using "mod_jk-1.2. 28-httpd-2.2.3.so".

2. Open the companion file provided with this Quick Start document and copy the mod_jk.conf, uriworkermap.properties and workers.properties files to the following location:

```
${APACHE_INSTALLED_LOC}\conf
```

If you do not have access to the companion file instructions on how to manually create these files can be found in Chapter 3.

3. Update the following lines in the workers.properties file to match the hostnames of your CA IT PAM nodes:

```
worker.node1.host=<DNS Name for Node 1>  
worker.node2.host==<DNS Name for Node 2>
```

In addition, if the ports or the nodename (node1 and node2) for the Orchestrators are changed this should also be updated in this file. However, this is not required for our basic example.

Note: You can add additional nodes to the cluster by following the instructions in "Customizing the workers.properties File" in Appendix A: Advanced Configurations of Apache HTTP Server and CA IT PAM.

4. Configure the Apache HTTP Server to load the mod_jk.so module by inserting the following lines along with the Include of conf files for other modules in the \${APACHE_INSTALLED_LOC}\conf\httpd.conf file:

```
#Load balancing module  
Include conf/mod-jk.conf
```

5. Restart the Apache HTTP Server and ensure that the server is up and running. The service can be restarted from the "Service Management Console" (services.msc) or by executing the following command from a command prompt:

```
C:> Net stop Apache2.2  
C:> Net start Apache2.2
```

At this point the Apache Server is set up as a loadbalancer for the CA IT PAM nodes that we are about to install.

Install CA EEM

CA IT PAM requires that an external directory (Security Server) is set up to manage users, groups and related security roles. Although Microsoft Active Directory (AD) and Sun One LDAP are supported we will be using CA EEM for our example as it is bundled with CA IT PAM.

To install CA EEM, do the following:

1. Ensure the CA IT PAM installation media is accessible from the Shared Server
2. Ensure that Java JRE 1.6 is installed or set the javahome parameter to "None" while launching the Install wizard (see note below).
3. Execute the following command from the \DVD2\EEM\ directory on the CA IT PAM media to launch the install

```
Win32_EEM_8.4.100.exe
```

Note: If you do not have Java JRE installed you need to specify this when launching the install wizard by using the following command:

```
D:\DVD2\EEM> Win32_EEM_8.4.100.exe -s -a /z"javahome=None; "
```

Note: The space between the; and the " is important.

4. Click Next on the Welcome Screen.
The License Agreement dialog appears.
5. Review the License Agreement, scroll down to the bottom of the text and select "I agree the terms in the license agreement". Then, click Next.
The Choose Destination Location dialog appears.
6. Click Next
7. Provide (and confirm) a password for EiamAdmin and click Next.
Important! It is critical that you remember this password!
8. If you are presented with the panel "Choose Java Home" browse to the Java Home directory (in other words, the directory that contains the "bin" sub-directory) and click Next.
9. The install will take a few minutes, Click Finish when you see the panel "CA Embedded Entitlements Manager Complete".

Create a CA IT PAM security certificate and security objects for CA EEM

Copy the CA IT PAM security configuration XML file ("ITPAM_eem.xml") from the \EEM sub-directory on the DVD2 of the CA IT PAM installation media to the \iTechnology sub-directory created when EEM was installed. By default this is C:\Program Files\CA\SharedComponents\iTechnology.

The default security certificate password is "itpamcertpass". This default may be changed by editing the value for the "password" attribute value in the following element in that document:

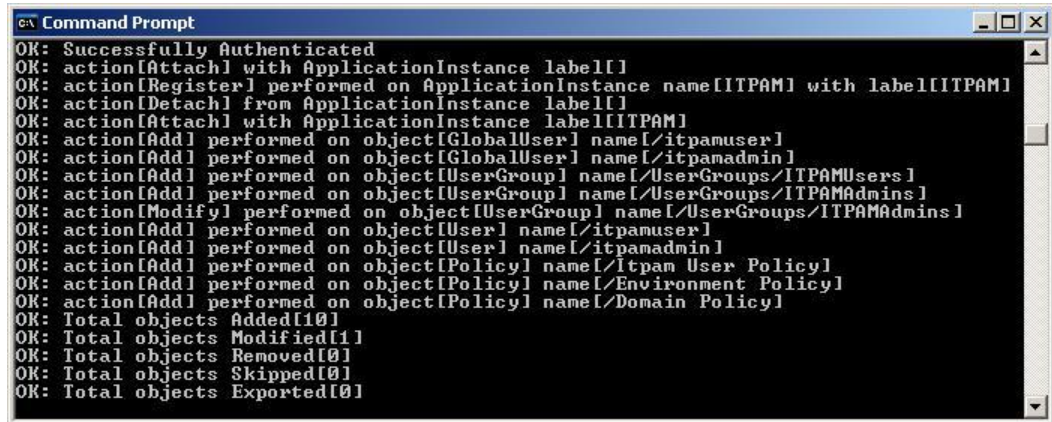
```
<register certfile="itpamcert.p12" password="itpamcertpass"/>
```



Next, open a command shell console window and navigate to the \iTechnology sub-directory. Confirm the "safex.exe" file exists. Execute the following command to create the IT PAM certificate file, EEM security groups and users:

```
safex.exe -h <hostname> -u EiamAdmin -p <Password> -f ITPAM_eem.xml
```

When executed correctly the output should be similar to the following:



```
CA Command Prompt
OK: Successfully Authenticated
OK: action[Attach] with ApplicationInstance label[]
OK: action[Register] performed on ApplicationInstance name[ITPAM] with label[ITPAM]
OK: action[Detach] from ApplicationInstance label[]
OK: action[Attach] with ApplicationInstance label[ITPAM]
OK: action[Add] performed on object[GlobalUser] name[/itpamuser]
OK: action[Add] performed on object[GlobalUser] name[/itpamadmin]
OK: action[Add] performed on object[UserGroup] name[/UserGroups/ITPAMUsers]
OK: action[Add] performed on object[UserGroup] name[/UserGroups/ITPAMAdmins]
OK: action[Modify] performed on object[UserGroup] name[/UserGroups/ITPAMAdmins]
OK: action[Add] performed on object[User] name[/itpamuser]
OK: action[Add] performed on object[User] name[/itpamadmin]
OK: action[Add] performed on object[Policy] name[/It pam User Policy]
OK: action[Add] performed on object[Policy] name[/Environment Policy]
OK: action[Add] performed on object[Policy] name[/Domain Policy]
OK: Total objects Added[10]
OK: Total objects Modified[1]
OK: Total objects Removed[0]
OK: Total objects Skipped[0]
OK: Total objects Exported[0]
```

This command will create the "itpamcert.p12" file in the \iTechnology subdirectory and populate required security objects for CA IT PAM in EEM. For security reasons, the CA IT PAM configuration XML file ("ITPAM_eem.xml") should be removed.

Set passwords for ITPAMAdmin and ITPAMUser

Proper execution of the "safex.exe" command creates not only the CA IT PAM security certificate file but the following required security groups and users in CA EEM:

- ITPAMAdmins group
- ITPAMUsers group
- ITPAMAdmin user
- ITPAMUser user

The default passwords for the "ITPAMAdmin" and "ITPAMUser" user IDs , are specified in the CA IT PAM XML configuration file for CA EEM, however, these values are – and must be – encrypted.

Since these passwords will be needed later, they must be reset using the CA EEM user interface. To do this:

1. Select Start, All Programs, CA,Embedded Entitlements Manager, EEM UI to launch the application
2. Provide the following information in the EEM login panel and then click "Log In".

Application:	ITPAM
User Name:	EiamAdmin
Password:	<password set when EEM was installed>

3. After successful login select the "Manage Identities" tab
4. Click Go in the "Search User" tile in the left hand pane



Two users ("itpamadmin" and "itpamuser") should be displayed in the "Users" tile under the "Users" folder. To reset the passwords:

1. Click on either user
2. In the right hand pane, scroll down in the "Global User Details" panel until the "Authentication" tile is visible
3. Select the "Reset Password" action from the "Authentication" tile
Two text boxes, "New Password" and "Confirm Password" should be displayed.
4. Enter the new password and confirm it
5. Click the "Save" button displayed at the bottom or top of the right hand pane
6. Repeat the process for the other user

The passwords for the "itpamadmin" and "itpamuser" should now be reset to known values. Log out and exit the CA EEM administrative UI.

Install CA IT PAM Domain Orchestrator on Primary Node

Install CA IT PAM Prerequisites

CA IT PAM requires the following three third party components:

- JBoss
- Hibernate
- JDBC

These components should be installed on each Orchestrator node; however, after the primary Domain Orchestrator is installed it is recommended that you deploy the remaining infrastructure components through the primary Domain Orchestrator.:

To install the required third party components on the primary Domain Orchestrator node do the following:

1. Verify that all prerequisites are met, including installation of the Java 1.6 JDK.
2. Launch "Third_Party_Installer_windows.exe" from DVD 1 of the CA IT PAM source media.
The Welcome dialog displays.
3. Click Next
The License Agreement dialog displays.
4. Review and click "I accept the agreement". Then click Next.
The Select Destination Directory dialog appears.
5. We recommend that you edit the default entry ("C:\Program Files\ITPAM") to read "C:\Program Files\CA\ITPAM".
The Prerequisites for CA IT PAM Installation dialog appears
6. Click Next
The JBoss Installation dialog will display.



JBoss Installation

The JBoss Installation dialog should display the path to the installation source “ZIP” file. By default this is: D:\ITPAM\DVD1\JBoss\jboss-4.0.3SP1.zip. If not, browse the correct path. Do not enable the “Use Domain” option. Click Next to proceed.

Note: Other versions of JBoss and Hibernate are not supported.

At this point the installer should begin unpacking and installing the required JBoss components to the \ITPAM directory. When completed, the “Hibernate Installation” panel will be displayed.

Hibernate Installation

On the “Hibernate Installation” Panel do the following:

1. Do not enable the “Use Domain” option
2. Browse to the correct source file if the default entry displayed (“D:\ITPAM\DVD1\hibernate\hibernate-3.0.jar”) is incorrect
3. Click Next

The installer will begin unpacking and installing the required Hibernate files. When complete the “JDBC Jars Installation” panel will be displayed.

JDBC Installation

1. On the “JDBC Jars Installation” panel do the following:
2. Click Add Files
3. Choose “MS SQL 2005” from the drop-down list.

The source location for the required JDBC JAR file should be displayed. By default this is D:\ITPAM\DVD1\drivers\sqljdbc.jar. If the source path is incorrect, click Browse to locate the “sqljdbc.jar”.

4. Click Next

The required file(s) will be copied to the proper location. The Tapi Jars Installation dialog will appear. However, since installation and configuration of telephony services is not required for this example, it will not be included in this document. Click Next to proceed.

Installation of Third Party Prerequisites Complete

Installation of third party prerequisites is now complete. Click Next on the resulting “Prerequisites for IT PAM installation” panel. The “Completing the 3rd Party Installer Setup Wizard” dialog will display.

Update the path to point to the source location for the CA IT PAM Installer. By default this is “D:\ITPAM\DVD2\”. Click “Finish” to launch the main CA IT PAM installation.

Install CA IT PAM Domain Orchestrator on Primary Node

The steps for installing the CA IT PAM Domain Orchestrator in a clustered environment are similar to installing a standalone Domain Orchestrator. The only difference is that, in a clustered environment, you will also need to identify the loadbalancer.



To install the Domain Orchestrator on the primary node in a clustered environment do the following:

1. If you are continuing directly from the 3rd party installer (see section above) you should see the "Welcome to the ITPAM domain Setup Wizard" dialog. If not you can start this part of the install by launching "CA_ITPAM_Domain_windows.exe" from the root of DVD 2.
2. Click Next on the "Welcome to the ITPAM Domain Setup Wizard" dialog.
The License Agreement dialog will display.
3. Review and select "I accept the agreement". Click Next to proceed.
4. If you are prompted to provide the Java Home Directory browse to the Java 1.6 JDK home directory (in other words, the directory that contains the "bin" sub-directory). Then, click Next.
5. If the panel "select JBoss Installation directory" is displayed ensure the path is correct (the same as you specified for the 3rd party components) and then click Next.
The ITPAM Domain Configuration dialog appears.
6. Ensure that the following configuration details for the loadbalancer are correct and that the loadbalancer is enabled:

Type of Server:	<i>New Server</i>
Load Balancer Host Name:	<i><dns name for the loadbalancer server></i>
Load Balancer Port Number:	<i>80</i>
Load Balancer Worker Node:	<i>node1</i>

Click Next to proceed.
7. For CA IT PAM r2.2: When prompted provide the name of your company and click Next.
The Set Certificate Password dialog appears.
8. Provide and confirm a password and click Next.
Note: This password is used to control access and encrypt password and other critical data. If you forget this password you will need to reinstall all Orchestrators.
The Select Start Menu folder dialog appears.
9. It is recommended you select "CA\ITPAM Domain". Click Next to proceed.
The General Properties dialog appears.
10. Accept the default values on this dialog unless it will result in a port conflict with other installed applications (TIP: Screen Capture the panel for future reference). Also, enable the "Install as Windows Service" and click Next.
11. Select EEM in the "Select Security Server Type" panel and click Next.
12. In the EEM Security Setting panel provide the following:
 - > In the "EEM Server" field enter the dns name for the EEM Server.
 - > In the "EEM Application Name" field enter ITPAM
 - > In the "EEM Certificate File" field enter the path to the itpamcert.p12 file that was created while running the safex command earlier (Default: "C:\Program Files\CA\SharedComponents\iTechnology\itpamcert.p12" on the EEM server).



- > In the "EEM Certificate Password" field enter the password that was specified in the "ITPAM_eem.xml" file when the itpamcert.p12 certificate was generated. If the XML file wasn't modified the default password is "itpamcertpass".

13. Before clicking Next, click the "Test EEM Settings" to verify the entries are correct.
14. Click OK in response to the "Performing a test...may take a few minutes to complete" message
15. When "Verify EEM settings" login dialog appears enter user name "itpamadmin" and the password previously set for the "itpamadmin" user in EEM and click OK.

Assuming all the correct values were entered a dialog indicating the following should appear:

Connect:	Ok
User provided belongs to User Group:	Ok
User is an Admin:	Yes

If necessary correct entries and re-test. Click Next when the test is successful.

16. In the "Database Settings" panel ensure the following are selected:

Type of Database:	MS SQL 2005 (or "MS SQL" for r2.2)
User Name:	sa
Password:	<password for the sa user>
Database Server:	<hostname for the database server>
Database Port:	<port to your SQL instance, default is 1433>
Repository Database:	itpam
Queues Database:	itpam
Driver Jar:	<path to sqljdbc.jar>

Note: Instead of "sa" you can use another user that has sufficient privileges to create a database, logins and grant privileges.

17. Create the "itpam" database by clicking the "Create Database" button. Click OK when the message "itpam has been created" appears.
18. Click the "Test Database Settings" button and click OK when the message "Test is successful" appears.
19. Click Next once you have successfully created and tested the database connectivity.
20. In the "Reporting Database Settings" panel check the "copy from main repository" checkbox. Verify that all information was correctly copied from the previous panel.
21. Click the "Test Database Settings" button and click OK when the message "Test is successful" appears.
22. Click Next once you have successfully tested the database connectivity.
23. In the "Additional Jars for Installation" dialog check the checkbox in front of the sqljdbc.jar and then click Next.
24. In the "Choose Connectors that needs to be installed/Updated" dialog select the connectors you need in your environment and then click Next.



Note: Installing the connectors does typically not require any user interaction, however there are some exceptions. Installation of the connectors is beyond the scope of this document.

25. When the Completing the ITPAM Domain Setup Wizard dialog appears, click Finish.

Launch the CA IT PAM Console

You need to start the CA ITPAM Server service before launching the CA IT PAM User Interface. This service is created by the installation. This can be done either through the service control manager (service.msc) or by running the following command from the command prompt:

```
C:> net start CAITPAMServer
```

Allow a few minutes after the service starts before attempting to launch the console. It is also recommended that you set the service to start up automatically on boot.

Then verify that the user interface is reachable through the loadbalancer, using the URL:

```
http://<itpam_loadbalancer >/itpam
```

Note: Java JRE or JDK 1.6 must properly installed and integrated with the browser on the remote system.

Login to the application using the "itpamadmin" user ID and corresponding password (see "Set passwords for ITPAMAdmin and ITPAMUser" on page 28). After successful login the "CA IT Process Automation Manager" user interface will open.

Blank Page Displayed

Depending on browser type, version and configuration a blank page may be displayed when you attempt to launch the management console. Typically, the cause of this is that the URL is not included in the current list of "trusted sites" for the browser. To resolve this issue, simply add the URL to the list of "trusted sites" for the browser and start the "ITPAM Domain Server" through the Service Console Manager

"Service Temporarily Unavailable" or "The Page cannot be displayed"

If the loadbalancer is misconfigured it might return a large number of error messages including, but not limited to, "Service Temporarily Unavailable" or "The Page cannot be displayed". A simple way to verify if the loadbalancer is at fault is to try to launch CA IT PAM directly from the primary node by using the URL:

```
http://<itpam_node1>:8080/itpam
```

If this works from the primary node but fails when launched through the loadbalancer fails then you may have a problem with the configuration of the loadbalancer (or the communication between the loadbalancer and the IT PAM node). Carefully review the section "Install the Apache HTTP Server" on page 25 for additional information.



Install Domain Orchestrator on Secondary Node

Installing the Domain Orchestrator Server on the Secondary node is very similar to installing the Domain Orchestrator on the Primary node except that most of the install panels will be pre-populated and read-only since it takes the information directly from the Primary Domain Server node.

The steps required to install Domain Server on the Secondary node are as follows:

1. Verify that all the prerequisites are met, including installation of the Java 1.6 JDK on the server that will function as the Secondary Node.
2. Open the CA IT PAM User Interface from the Secondary Node by launching the following URL:
`http://<itpam_loadbalancer>/itpam`
3. Enter the user "itpamadmin" and the password you created earlier (in the "Set passwords for ITPAMAdmin and ITPAMUser" on page 28) and click on Log In.
4. Select the "Installation" section in the left hand pane of the "CA IT Process Automation Manager" user interface.



5. Click on the "Install Secondary Domain Server" subsection below "Installation" in the left hand pane.
6. In the corresponding right hand pane click on the "Go >>>" button. This will launch the Install Wizard for CA IT PAM and its related 3rd party components.
7. If an "Authentication Required" message is displayed cancel the installation and ensure that your browser or Java environment is set up so that bypasses any proxy for the loadbalancer node. For information on how to do this consult the documentation or online help provided with your browser.
8. If a panel indicating "The application's digital signature cannot be verified." is displayed, check "Always trust content from this publisher" and click on Run.
The Welcome dialog displays.
9. Click Next
The License Agreement dialog displays.
10. Review the license agreement, select "I accept the agreement" and click Next.
11. Update the Destination Directory to "C:\Program Files\CA\ITPAM" and click Next

The Prerequisites for ITPAM Installation dialog displays.

12. Click Next

The JBoss Installation dialog displays.

13. Ensure the checkbox for "Use Domain" is checked and click Next.

The Hibernate Installation dialog displays.

14. Ensure the checkbox for "Use Domain" is checked and click Next

The Prerequisites for ITPAM installation dialog displays.

15. Click Next.

The Completing the 3rd Party Installer Setup wizard dialog displays.

16. Ensure the checkbox for "Use Domain" is checked and click Finish.

Note: It might take a few minutes to copy the IT PAM Installer to the local server.

17. When the Welcome to the ITPAM Domain Setup wizard dialog displays, click Next

The License Agreement dialog displays.

18. Review the license agreement, select "I accept the agreement" and click Next.

19. If you are presented with a panel asking for the "Java Home Directory" browse to the Java 1.6 JDK home directory (i.e. the directory that contains the "bin" sub-directory) and then click Next.

20. In the "ITPAM Domain Configuration Screen", ensure the configuration for the loadbalancer is enabled and correct:

Type of Server:	<i>Additional Cluster Node</i>
Load Balancer Host Name:	<i><dns name for server loadbalancer></i>
Load Balancer Port Number:	<i>80</i>
Load Balancer Worker Node:	<i>node2</i>

When this is correct click Next.

The Set Certificate Password dialog displays.

21. Provide the password that was created for the Primary Domain Orchestrator and click Next.

The Select Start Menu Folder dialog displays.

22. We recommend that you specify "CA\ITPAM Domain" and click Next.

23. On the "General Properties" panel accept the default unless it will result in a port conflict with other installed applications (TIP: Screen Capture the panel for future reference). Also, enable the "Install as Windows Service" and click the "Next" button.

24. Click Next in the "Select Security Server Type" panel.

25. Click Next in the "EEM Security Setting panel"

26. Click Next in the "Database Settings" panel



27. Click Next in the “Reporting Database Settings” panel

28. When the panel “Completing the ITPAM Domain Setup Wizard” appears click Finish.

You need to start the CA IT PAM Server service before attempting to use this new CA IT PAM node. This service is created by the installation process and can be started either through the service control manager (service.msc) or by running the following command from the command prompt:

```
C:> net start CAITPAMServer
```

Allow a few minutes after the service starts before attempting to launch the console. It is also recommended that you set the service to start up automatically on boot.

Verify Cluster Functionality

The final step is to verify that the cluster configuration works as expected in the event one of the nodes becomes unavailable. This section describes a basic test that can be used to:

- Perform a basic test of the loadbalancer and cluster nodes
- Test failover of CA IT PAM while a process is being edited
- Test failover of CA IT PAM with an active workflow process

Perform a Basic Test of the Loadbalancer and Cluster Nodes

This first basic test is to verify that you can reach CA IT PAM. When the cluster configuration is working correctly, you should be able to access an active CA IT PAM node regardless of which nodes are available. This should occur transparently. In other words, there is no need for you to specify a particular node.

First, ensure that all servers and services (APACHE2.2 on the load balancer node and CAITPAMServer on both ITPAM nodes) are up and running. Next, launch and login (using itpamadmin) to CA IT PAM through the load balancer node:

<http://<itpam loadbalancer>/itpam>

Next, close the browser and shut down one of the IT PAM nodes (either by shutting the complete server or by shutting down the CAITPAMServer service) and launch a new browser through the loadbalancer. After you have logged in and verified basic CA IT PAM functionality on this node you can bring up this server again and then do the same test for the other IT PAM node.

If you have any problems connecting to one or both nodes it is recommended that the first step in the troubleshooting is to verify if you can connect directly with the individual CA IT PAM nodes. This can be done using the URL below:

http://<itpam_node>:8080/itpam

If this direct connection works you need to troubleshoot the loadbalancer setup. A few basic steps are:

- Carefully verify the install and configuration steps listed in the section “Install the Apache HTTP Server”.
- Verify that you haven’t modified any of the ports for the loadbalancer or CA IT PAM
- Verify that you don’t have a firewall blocking the traffic between the loadbalancer and the CA IT PAM nodes
- Finally review the documentation for the Apache HTTP Server and the mod_jk module on the <http://apache.org> site.

Test Failover of CA IT PAM while a Process is Being Edited

The next test of the cluster is to verify that active CA IT PAM sessions can failover transparently between the nodes. Any easy method to verify this is to:

1. Shut down the secondary CA IT PAM node (or stop the CAITPAMServer service) and open the CA IT PAM interface through the loadbalancer.

http://<itpam_loadbalancer>/itpam

Note: Do not use a web browser on any of the CA IT PAM nodes for this. Either use a separate client or a webbrowser hosted on the loadbalancer.

Since Node 2 is down you know that you are connected to Node 1.

2. Log in to CA IT PAM using the “itpamadmin” user ID and the password you associated with this user (see “Set passwords for ITPAMAdmin and ITPAMUser” on page 28).
3. Open the CA IT PAM Client by clicking on the link “ITPAM Client” in the upper right corner of the webbrowser.

Note: If this is the first time the UI is opened it will clearly state what node it is downloading the interface from (In our example it should be <http://node1:8080>)

4. To prepare for the failover this is a good time to bring up CA IT PAM on node 2. Start the server and verify that the “CAITPAMServer” service is started. At this point, both nodes are active, however we know for sure we logged into node 1.
5. Open the “Library Browser” through the menu item “File/Open Library Browser”
6. Create a folder for the cluster test called “ClusterTest” by right clicking on root “/” folder and select “New Object / Folder”. This creates a new folder in the tree; rename this folder to “ClusterTest”.
7. Create a test process by right clicking on newly created folder and select “New Object / Process”. This creates a new process in the Right hand pane; rename this Process to “ClusterTestProcess1”.
8. Open this process by double clicking on it; this will open a new blank window for this process.



9. Drag and drop a "Start" operator from the "Common console" in the left hand pane into this new workspace in the right hand pane, and then save the new process (click on the save icon or through the "File/Save" menu option).

10. In the same way drag and drop a "Normal Stop" operator from the common console.

At this point we have logged into node 1 and started to edit a process in this environment. We saved a small piece of it but then continued to edit the process without saving it. This is a great time to forcefully take down node 1 and verify that the cluster truly works.

11. Log into CA IT PAM Node 1 and shut down the system, or, if you prefer, you can just shut down the service by running the following command:

```
C:> net stop CAITPAMServer
```

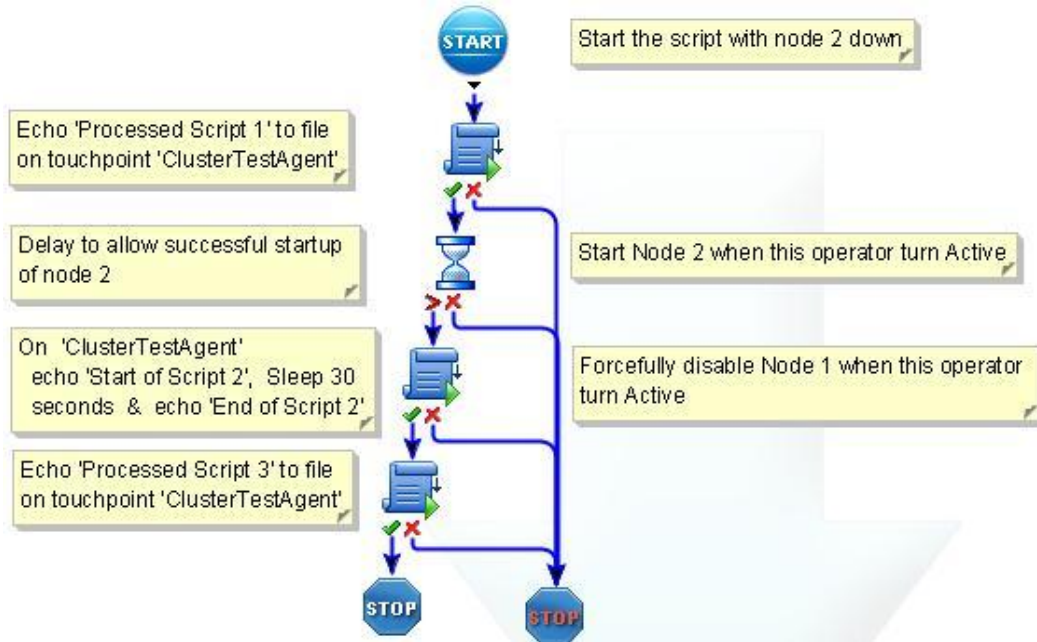
12. When the node 1 is confirmed down go back to the "CA IT PAM Client" where you are editing the process and drag and drop out an "Abnormal Stop" operator from the "Common Console".
13. Save the process by clicking on the save icon (or by using the "File/Save" menu), since node 1 is down this is now handled by node 2.

What you have done at this point is to ensure that the failover of the CA IT PAM works transparently when one of the nodes goes down.

Test Failover of CA IT PAM with an Active Workflow Process

The final test of your newly created cluster is to ensure that it can correctly handle processes that are active and in flight in the event there is a catastrophic failure of the Orchestrator that is responsible for it. Following is a high level overview of the approach for this test:

1. Install a CA IT PAM Agent on a server (not on any of the Orchestrator nodes) and configure a Touchpoint for this Agent. A dedicated CA IT PAM Agent test server or the server for shared components (Load balancer/EEM) can be used for this.
2. Finish the "ClusterTestProcess1" so that it includes the following 4 operators:



- > Operator 1: Script that writes "Processed Script 1" to a file on the Agent server
- > Operator 2: Sleep function that gives you time to start up node two (allowing you to start on node one and then give node 2 enough time to start up)
- > Operator 3: Script that writes "Start of Script 2", sleep for 30 seconds to allow you to disable node 1 and then finally writes an 'End of Script 2' message into the log file on the Agent server.
- > Operator 4: Script that writes 'Processed script 3' message. This ensures that the handover to any additional steps works as expected.

This workflow is included with the name "ClusterTestProcess1_CompFileVer" in this document's companion file.

3. Ensure that the IT PAM service on node 2 is stopped and then start the process. This ensures the process starts on node 1.
4. When the process reaches Operator 2 start up IT PAM service on node 2
5. When the process reach Operator 3 abruptly shut down node 1 where the process was started (ensuring it can communicate and hand over the status to node 2)
6. Verify that the process is successfully executed (with all expected messages on the agent server)

Following are detailed instructions on how to accomplish these tasks.

Installing the CA IT PAM Agent

1. Verify that Java 1.5 JRE or above is installed on the server on which you will be installing the Agent. . This can be the load balancer node or a new server.
2. Open the CA IT PAM User Interface from the Agent Server by launching the following URL:
`http://<itpam_loadbalancer>/itpam`
3. Enter the user "itpamadmin" and the password you created earlier (in the "Set passwords for ITPAMAdmin and ITPAMUser" on page 28) and click Log In.
4. Select the "Installation" section in the left hand pane of the "CA IT Process Automation Manager" user interface.



5. Click on the "Agent Installation" subsection below "Installation" in the left hand pane.
6. In the corresponding right hand pane click on the "Go >>>" button. This will launch the Install Wizard for the IT PAM Agent.
7. If a "File Download - Security Warning" message is displayed click on "Run"
8. If a panel indicating "The publisher could not be verified..." click on Run.
The Welcome dialog displays.
9. Click Next
The License Agreement dialog displays.
10. Review the license agreement, select "I accept the agreement" and click Next.
11. If you are presented with a panel asking for the "root directory of the JRE installation" browse to the Java JRE home directory (i.e., the directory that contains the \bin sub-directory) and then click Next.
12. Update the Destination Directory to "C:\Program Files\CA\ITPAM Agent" and click Next
The Select Start Menu Folder dialog displays.
13. We recommend that you specify "CA\ITPAM Agent" and click Next.
14. In the ITPAM Agent Domain URL provide the URL for your domain

<http://<loadbalancer>/>

15. On the "General Properties" panel accept the defaults unless it will result in a port conflict with other installed applications (TIP: Screen Capture the panel for future reference). Also, enable the "Install as Windows Service" and "Start Agent after Installation" checkboxes, finally click the "Next" button.
16. When the panel "Completing the ITPAM Agent Setup Wizard" appears click Finish.
17. Use the Service Manager (services.msc) to verify that the Agent Service "CA IT PAM Agent" is successfully started.

Configure a CA IT PAM Touchpoint Referencing the Agent

1. Open the IT PAM interface through the loadbalancer.
http://<itpam_loadbalancer>/itpam
2. Log in to CA IT PAM using the "itpamadmin" user ID and the password you associated with this user (see "Set passwords for ITPAMAdmin and ITPAMUser" on page 28).
3. Open the "IT PAM Client" by clicking on the link "ITPAM Client" in the upper right corner of the webbrowser.
4. Select the tab "Configuration Browser" if it is open. Otherwise open it through the menu item "File/Open Configuration Browser"
5. Lock the environment by right clicking on the node "Default Environment" (in the left hand pane) and selecting the "Lock" option.
6. Once again, right click on the "Default Environment" and select the menu option "Add Touchpoint"
7. In the Window "Add Touchpoint: Default Environment" specify the Touchpoint name "ClusterTestAgent" and select the Agent you just installed from the list of agents. Click on the Save button in this window.
This brings you back to the "Configuration Browser".
8. Save the updated environment by clicking on the Save icon (or by using the "File/Save" menu selection)
9. Unlock the environment by right clicking on the node "Default Environment" and selecting the "Unlock" option.

Finalize "ClusterTestProcess1" process to support this test

Once you have installed an agent and created a related Touchpoint you need to create a simple test process that can be executed while testing the failover functionality. A sample process is provided in the companion file to this document. To import the file from the companion file do the following:

1. Extract the "ClusterTestProcess1_CompFileVer.xml" file from the companion file.
2. Open the "Library Browser" through the menu item "File/Open Library Browser"
3. Right click on the "ClusterTest" folder and select "Import"
4. Navigate to and select the "ClusterTestProcess1_CompFileVer.xml" that you just extracted. Click Open.



5. Check the "Set imported version as current" option and click OK

Once you have successfully imported the process proceed to the "Execute the Failover Test of the Active IT PAM Process" section on page 44.

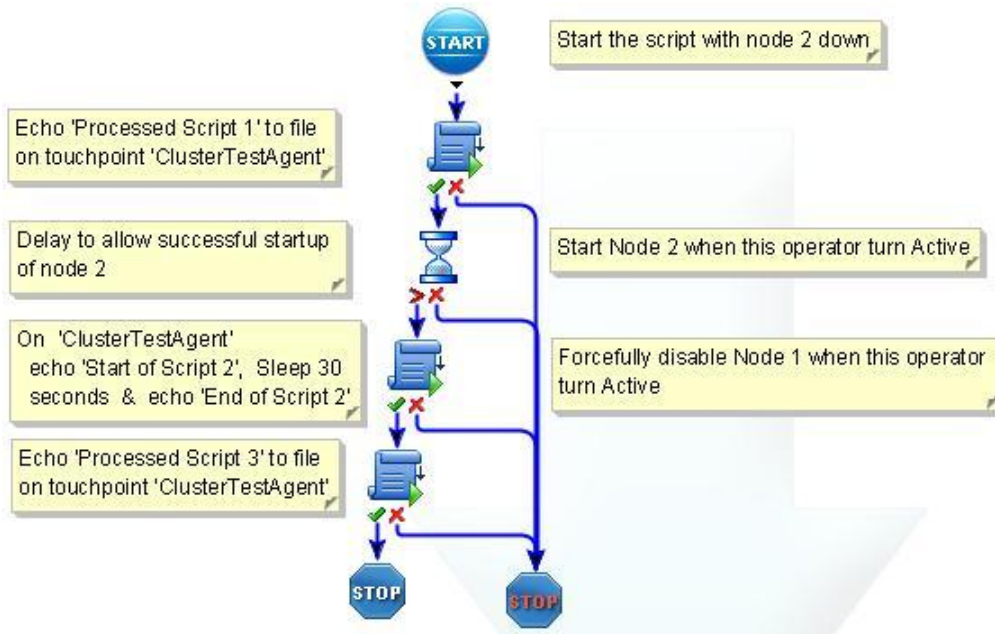
You can also create the process manually using the following steps:

1. Open the "Library Browser" through the menu item "File/Open Library Browser"
2. Double-click on the "ClusterTestProcess1" process that you created in the "ClusterTest" folder during the "Test Failover of CA IT PAM while a Process is Being Edited" section on page 37. This process should, at this point, include a "Start", a "Normal Stop" and an "Abnormal Stop" operator.
3. Drag and drop four new operators into this process:
 - > Operator 1: A "Start Script" operator from the "Process Module"
 - > Operator 2: A "Delay" operator from the "Date-Time Module"
 - > Operators 3 & 4: Two "Start Script" operators from the "Process Module"

If a module is not directly visible in the left hand pane, it can be reached by clicking on the double arrow ">>" next to "Edit User Preferences".

4. On each of these 3 "Start Script" operators right click on the corresponding icons and select "Completed" and "Failed". This will add two ports at the bottom of the operator.
5. Right click on the "Delay" operator and select "After" and "Failed".
6. Click and drag the small "port" (an upside down triangle) from the start operator to the top of the Operator 1 icon. This will connect the two with an arrow pointing from the Start operator to Operator 1
7. Using the same procedure connect the following:
 - > The completed port ("Green Checkmark") from Operator 1 to Operator 2
 - > The after port ("red >") from Operator 2 to Operator 3
 - > The completed port ("Green Checkmark") from Operator 3 to Operator 4
 - > The completed port ("Green Checkmark") from Operator 4 to the Normal Stop operator.
 - > The failed port ("red x") from all operators to the Abnormal Stop

The end result should look like this (with the exception of the yellow notes stickers):



8. Doubleclick on Operator 1 to open its "properties" for configuration.
 - > Go to the script section and click on "Inline Scripts". Enter the following scripts and then click OK.

```
echo Processed Script 1 >> C:\ClusterTest.log
```
 - > In the "Execution Settings" click Select. Locate and select the "ClusterTestAgent" Touchpoint and click OK
 - > Click Apply at the bottom of the properties pane.
9. Open the properties pane of Operator 2 and enter 60 in the "Seconds" field (in the "Delay" section) and then click "Apply".

Note: It is critical that the IT PAM Service starts and stabilizes on node two before node one is taken down. Therefore, if the Orchestrators (especially node 2) are running on slow servers you might consider making this delay longer than 60 seconds.
10. Modify Operator 3 in the same way as you did with Operator 1, except use the following script:

```
echo Start of Script 2 >> C:\ClusterTest.log  
ping localhost -n 30 > NUL  
echo End of Script 2 >> C:\ClusterTest.log
```

Note: The ping command is used as a sleep to wait approximately 30 seconds
11. Modify Operator 4 in the same way as you did with Operators 1 and 3, except use the following script:

```
echo Processed Script 3 >> C:\ClusterTest.log
```
12. Save the project by clicking on the save icon at the top of the CA IT Process Automation Manager User Interface (or by using the "File/Save" menu selection).

At this point, all of the test preparations are done. The next step is to put it all together and verify the failover functionality.

Execute the Failover Test of the Active IT PAM Process

Before you start this last step you need to decide how abruptly you would like to abort the Orchestrator on node 1. One useful and relatively safe simulation is if the NIC or network is disabled. Although you can also power off the server or shut down the OS or ITPAM service, these options are not recommended since powering off the system in extreme cases can cause corruption and shutting down the OS or ITPAM service would allow the service to handover a status message to the remaining nodes. Once you have selected your approach, you can start the test.

1. Shut down the IT PAM Service on node 2

```
C:> net stop CAITPAMServer
```

2. Prepare - but do not execute - a quick method to start node 2. For example, a suggested method is to open a command window on node 2 and type the following command:

```
C:> net start CAITPAMServer
```

Do not hit Enter to execute this command yet!

3. Prepare – but do not execute - a quick method to disconnect the IT PAM service on node 1 from the environment. If you have physical access to the server, a recommended method is to disconnect the network cable.
4. Open the “ClusterTestProcess1” (or “ClusterTestProcess1_CompFileVer” if you imported the predefined process) that you just created if it is not already opened.
5. Start the process by clicking on the green “play” button (or by using the “Control/Start” menu selection)
6. Click “Yes” in the “Monitor Process Instance” window

A new window indicating how the process moves along is opened. An Orchestrator with a purple tint indicates the currently running operator and a green icon indicates a successfully executed operator.

7. As soon as Operator 2 (the Delay operator) becomes active (turns purple) you should start the IT PAM Service on node 2 (see Step 2 above)
8. When Operator 3 becomes active disconnect the CA IT PAM service on node 1 (see Step 3 above).

Since we know that the process started on node 1 this will force a failover to the newly started node 2.

9. Wait and ensure that all operators, except the “Start” and “Abnormal Stop”, turn green. When the “Normal Stop” icon turns green the process has successfully completed.
10. On the Agent Server verify that the file C:\ClusterTest.log was created and include the following text:

```
Processed Script 1  
Start of Script 2  
End of Script 2  
Processed Script 3
```

This was the final test of the cluster. At this point you have proven that the CA IT PAM Orchestrator can successfully connect work in a clustered environment. This includes both high availability of user interface used to edit processes as well as the ability to manage currently active workflow during a catastrophic failure.

Appendix A: Advanced Configurations of Apache HTTP Server and CA IT PAM

This chapter provides information about the contents of the companion file provided with this paper.

Mod_jk.so Module

The Mod_jk.so module is used to connect Tomcat to the Apache HTTP server and it can be used to configure the Apache HTTP server as load balancer for IT PAM. The document includes a companion file with files that needs to be added and configured to correctly set up the loadbalancing environment. The three included files are:

- mod-jk.conf file
- workers.properties file
- uriworkermap.properties file

The section below outline the content of these files and can be used if you have misplaced the companion file.

Contents of the mod-jk.conf file

The Companion File to this paper includes a mod-jk.conf file which is used to configure the mod_jk.so module for the Apache HTTP Server. As described earlier in this paper this file should be copied to the following location:

```
${APACHE_INSTALLED_LOC}\conf
```

If you do not have access to the companion file you can create it manually by following the steps below:

1. Create a new text file called mod-jk.conf
2. Add the following information to this newly created file:

```
# Load mod_jk module
# Specify the filename of the mod_jk lib
LoadModule jk_module modules/mod_jk.so

# Where to find workers.properties
JkWorkersFile conf/workers.properties

# Where to put jk logs
JkLogFile logs/mod_jk.log

# Set the jk log level [debug/error/info]
JkLogLevel error

# Select the log format
JkLogStampFormat "[%a %b %d %H:%M:%S %Y]"
```



```
# JkOptions indicates to send SSK KEY SIZE
JkOptions +ForwardKeySize +ForwardURICompat -ForwardDirectories

# JkRequestLogFormat
JkRequestLogFormat "%w %V %T"

# You can use external file for mount points.
# It will be checked for updates each 60 seconds.
# The format of the file is: /url=worker
# /examples/*=loadbalancer
JkMountFile conf/uriworkemap.properties

# Add shared memory.
# This directive is present with 1.2.10 and
# later versions of mod_jk, and is needed for
# for load balancing to work properly
# JkShmFile logs/jk.shm

# Add jkstatus for managing runtime data
<Location /jkstatus/>
    JkMount status
    Order deny,allow
    Allow from all
</Location>
```

Contents of the workers.properties file

Also included in the Companion File is a workers.properties file which is used to configure the mod_jk.so module for the Apache HTTP Server. As described earlier in this paper this file should be copied to the following location:

```
${APACHE_INSTALLED_LOC}\conf
```

If you do not have access to the companion file you can create the workers.properties file manually by following the steps below:

1. Create a new text file called workers.properties under the \${APACHE_INSTALLED_LOC}\conf folder
2. Add the following information to this newly created file:

```
#Define list of workers that will be used for mapping requests
worker.list=loadbalancer, status, Primaryloadbalancer

# Define node1
# modify the host as your host IP or DNS name.
worker.node1.port=8009
worker.node1.host=<DNS Name for Node 1>
worker.node1.type=ajp13
worker.node1.lbfactor=1

# Define Node2
# modify the host as your host IP or DNS name.
worker.node2.port=8009
worker.node2.host=<DNS Name for Node 2>
worker.node2.type=ajp13
worker.node2.lbfactor=1

# Load-balancing behaviour
worker.loadbalancer.type=lb
worker.loadbalancer.balance_workers=node1,node2
worker.loadbalancer.sticky_session=1
worker.loadbalancer.retries=1
```



```
# Mirroring Load-balancing behaviour
worker.Primaryloadbalancer.type=lb
worker.Primaryloadbalancer.balance_workers=node1,node2
worker.Primaryloadbalancer.sticky_session=1
worker.Primaryloadbalancer.retries=1

# Status worker for managing load balancer
worker.status.type=status
```

Customizing the workers.properties File

The workers.properties file needs to be customized to fit your environment. In a basic two node cluster without any customized ports you only need to provide the DNS name for the two nodes.

The most common configurations done in this file are as follows:

- Specifying the DNS name for the CA IT PAM nodes. This is done by modifying the following rows:

```
worker.node1.host=<DNS Name for Node 1>
worker.node2.host==<DNS Name for Node 2>
```

To something like

```
worker.node1.host=itpam1
```

- Adding additional nodes. This require you to add a section similar to the following:

```
# Define nodeX
# modify the host as your host IP or DNS name.
worker.nodeX.port=8009
worker.nodeX.host=<DNS Name for Node X>
worker.nodeX.type=ajp13
worker.nodeX.lbfactor=1
```

As well as adding references to this new node X in the following two lines...

```
worker.loadbalancer.balance_workers=node1,node2,nodeX
worker.Primaryloadbalancer.balance_workers=node1,node2,nodeX
```

- Change the naming convention for the “Load Balancer Worker Node”. In our example they are called node1, node2 etc. This naming standard can be changed by modifying all references to these names in this file as well as in the Load Balancer panel in the CA IT PAM install wizard.
- Modifying the port ajp uses to communicate with CA IT PAM. By default this is port 8009. To change this, modify the port number in each node’s row as follows:

```
worker.nodeX.port=8009
```

Also modify this in each Orchestrator’s OasisConfig.properties file. This file can be found in the following location:

```
<ITPAM_DIR>/server/c2o/.config/OasisConfig.properties
```

In this file you should make sure the following row specifies the same port number as specified in the workers.properties file:

```
tomcat.connector.ajp.port=8009
```



Contents of uriworkermap.properties file for a Domain Orchestrator

The Companion File to this paper also includes an uriworkermap.properties file which is used to configure the mod_jk.so module for the Apache HTTP Server. As described earlier in this paper this file should be copied to the following location:

```
${APACHE_INSTALLED_LOC}\conf
```

If you do not have access to the companion file you can create it manually by following the steps below:

1. Create a new text file called uriworkermap.properties under the
\${APACHE_INSTALLED_LOC}\conf folder
2. Add the following information to this newly created file:

Note: This is unique to the Domain Orchestrator. The contents of this file for regular (non-Domain) Orchestrators is different.

```
#In uriworkermap.properties for Domain clustering make following entries:
# Mount the Servlet context to the ajp13 worker
/jmx-console=loadbalancer
/jmx-console/*=loadbalancer
/web-console=loadbalancer
/web-console/*=loadbalancer

# Mount your applications
/itpam/*=loadbalancer
/itpam=loadbalancer

# All the mirroring requests will go to the primary domain server
/itpam/MirroringRequestProcessor=Primaryloadbalancer
/itpam/MirroringRequestProcessor/*=Primaryloadbalancer
/c2orepository=loadbalancer
/c2orepository/*=loadbalancer
/c2orepository/media=Primaryloadbalancer
/c2orepository/media/*=Primaryloadbalancer
/c2orepository/thirdParty=Primaryloadbalancer
/c2orepository/thirdParty/*=Primaryloadbalancer

#Request to upload report and resource will go to primary domain server
/itpam/clientproxy/c2oresourceaction=Primaryloadbalancer
/itpam/clientproxy/c2oreportaction=Primaryloadbalancer

#All the mirroring request will go to primary server
/mirroringrepository=Primaryloadbalancer
/mirroringrepository/*=Primaryloadbalancer

#Agent start up request will go to primary domain server
/itpam/StartAgent=Primaryloadbalancer
/itpam/StartAgent/*=Primaryloadbalancer

#Gwt requests which are specific to primary domain server like Manage Version
#and Reporting will go to primary domain server only
/itpam/OasisPrimary=Primaryloadbalancer

#Installation request will go to primary domain server
/c2orepository/htmlFile/installation/*=loadbalancer

#All the Secondary domain setting setting request will go to Primary domain servers
/itpam/ServerConfigurationRequestServlet=Primaryloadbalancer
```



```
#Agent installation request will go to load balancer
/itpam/AgentConfigurationRequestServlet=loadbalancer

#All the reports will be uploaded to the primary domain server only. So we will have
#to map the URL which will be used to open the report to primary domain server.
/birt/*=Primaryloadbalancer

#The request for the oasis client should go to load balancer
/itpam/JNLRequestProcessor=loadbalancer
/itpam/JNLRequestProcessor/*=loadbalancer

#The request for the third party installation should go to primary server
/itpam/JNLRequestProcessor/installation=Primaryloadbalancer
```

Contents of uriworkemap.properties for non-Domain Orchestrators

The information below is not required for the example provided in this paper. However, it is useful information for doing similar clustered setup for non-Domain Orchestrators.

Note: This information only applies to basic (non-Domain) clustered Orchestrators. See the previous section for examples applicable to Domain Orchestrators.

The contents of the uriworkemap.properties when loadbalancing Orchestrators other than the Domain Orchestrator are as follows:

```
#In uriworkemap.properties for Normal server clustering make following entries:
/jmx-console=loadbalancer
/jmx-console/*=loadbalancer
/web-console=loadbalancer
/web-console/*=loadbalancer

# Mount your applications
/itpam/*=loadbalancer
/itpam=loadbalancer

#Help request can go to any server, hence mapping request to loadbalancer
/c2orepository/oasisHelp=loadbalancer
/c2orepository/oasisHelp/*=loadbalancer
/c2orepository/htmlFile/aboutUs/*=loadbalancer
/c2orepository/htmlFile/language/*=loadbalancer

#We want secondary server configurations settings to come from primary server only.
/itpam/ServerConfigurationRequestServlet=Primaryloadbalancer

#Mirroring among the normal cluster server
/itpam/MirroringRequestProcessor=Primaryloadbalancer
/itpam/MirroringRequestProcessor/*=Primaryloadbalancer
/c2orepository/*=loadbalancer

#all the cluster related files download request will go to primary
/mirroringrepository=Primaryloadbalancer
/mirroringrepository/*=Primaryloadbalancer
```