GROUP POLICY CENTRAL

Information about Group Policy for IT Administrators

Group Policy Design Guidelines – Part 2

Posted by Alan Burchill on 27 July 2010, 7:00 pm

In my previous article In this article Best Practice:Active Directory Structure Guidelines – Part 1 I spoke about some of the guidelines I personally use when developing an Active Directory OU structure. In this next part I will discuss some guidelines I use when designing a Group Policy Object infrastructure.

Ideally you should make the the Active Directory OU and GPO design decision together to best ensure that you have the most efficient design possible. However if you have an existing OU structure designed a lot of these guidelines can still be applied to most existing environments.

As in Part 1 these are simply guidelines that I use and should not be taken as hard an fast rules. I quite often finding myself having to break these rules due to real world conflicts or just because one rule might conflict with the other rule. If you do find your self in a situation where you are not sure which path to take try to chose the option that will result in the least administrative effort in the long term.

Active Directory Group Policy Design Guidelines

Keep the GPO's name consistent with the OU names

When naming the GPO try to keep the name of the policy the same as the concatenated name of all the OU's to where the group policy object is applied. Having the fully concatenated name will make it intently know what that policy is applied when just looking at the GPO name. This is very handy to know when looking at a Group Policy Results report which only gives you the name of the GPO without the linked OU details.

Bad Example "Workstations"	Good Example "Sydney Workstations"
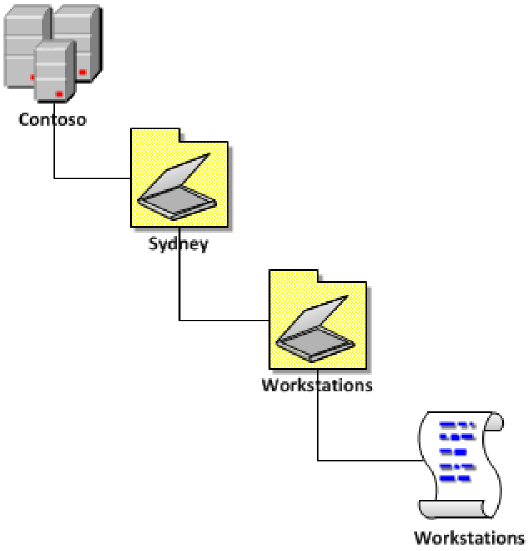	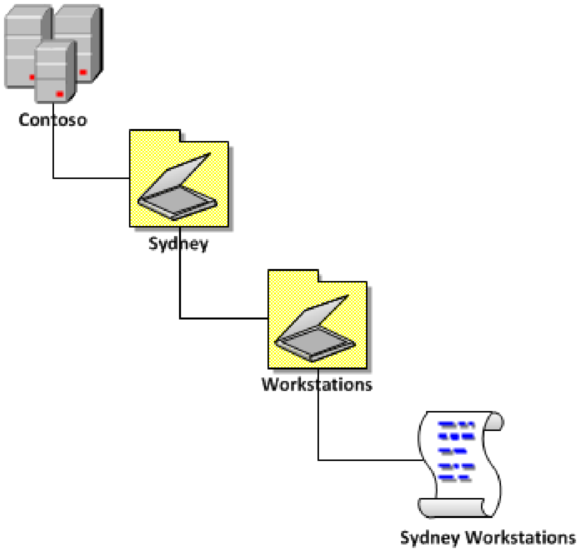

In keeping with having names consistent this also means you should adhere to the same naming conventions as mentioned in Part 1 with the OU's (i.e. "Keep it short", "Be Intuitive" & "Most to least signification from left to right"… So in saying that please read the next guideline…

REFERENCES

TechNet: Establishing Group Policy Operational Guidelines

> Define a meaningful naming convention for GPOs that clearly identifies the purpose of each GPO

Don't use the work "POLICY" or "GPO" in the GPO name

Nothing annoys me more to see a group policy called "Workstations Policy" or "Workstation GPO"…. I KNOW ITS A POLICY!!!! I AM LOOKING AT IT IN THE GROUP POLICY MANAGEMENT CONSOLE. Please drop the work "policy" or "GPO" from the name of the Group Policy object as you are simple adding more characters to what might already be a long name only for the sake of pointing out the obvious.

I also realise that the two GPO's that come with AD are called "Default Domain Policy" and "Default Domain Controller Policy" which goes against this rule…

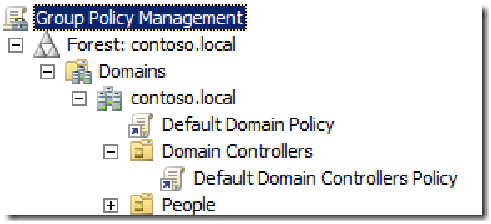

Remember at the start of part 1 how rules were meant to be broken…   So I do NOT recommend that you rename these polices there is just to much risk and confusion that doing this might cause. But this would have to be the only exception to this rule that I would be happy to let though…

Treat your terminal servers like workstations

Terminal Servers (now known as Remote Desktop Services) are essentially a multi-user workstation and as such should be treated more as a workstation than a server. Ideally you should configure you Terminal Server to be as close as possible as your workstations to provide your users with a consistent experience. The best way to make sure the configuration is consistent is to apply the same policy settings to the Terminal Serves as your workstations.

That being said don't apply the same computer Group Policy Object to the Terminal Servers if for no other reason than it helps reduce the risk of making a change to a workstation that could affect the stability of the servers (e.g. Automatic Update reboot schedule). Therefore you will need to maintain some level of manually synchronisation between you default workstation and terminal server policy.

Unlike computer GPO's it far more acceptable to apply the same user GPO's to your users when logging on to the Terminal Server as the GPO are applied to the User Object rather than the computer account. Using the same policy means that any changes made to the user policies will automatically apply to terminal servers without the administrative overhead of making duplicate updates when there are policy changes. If you have any user configuration that you want to configure that is specific to the terminal servers (e.g. disable adding PST file) then you can override this policy using the Group Policy Loopback option on the computer GPO you apply to the Terminal Server. This is another   reason why you would want to have a separate computer GPO as it allow you to apply specific Terminal Server user settings via a loopback policy.

For more information on troubleshooting Loop back policies check out Loopback Policy Processing Debug Series | CB5 Blog and Aaron Parker's StealthPuppy blog.

REFERENCE

TechNet: Using Loopback Processing to Configure User Settings

> The **User Group Policyloopback processing mode** policy setting is an advanced option that is intended to keep the configuration of the computer the same regardless of who logs on. This option is appropriate in certain closely managed environments, such as servers, terminal servers, classrooms, public kiosks, and reception areas.

New GPO's only when scope is different

I have seen some organisations apply many Group Policy Objects (GPO's) to the same OU. There are a number of reason why you might want to do this however you should really consider why you want spawn another GPO as each one will add about 5mb to you Active Directory SYSVOL. But if you start creating lots of GPO objects then you can quickly blow out your the size and performance of your SYSVOL. This is not such a problem if you have upgraded to a DFS-R SYSVOL replication or you have configured a Group Policy Central Store for your Windows Vista and later computers but its still good practice to keep the number of GPO's as low as possible.

Monolithic vs. Functional GPOs

Now that I have just told you that you should load up your GPO's with lots of setting rather than having lots and lots of separate GPO's Mike Kline has referred me to the this great article Best Practice for Optimizing Group Policy Performance by Darren Mar-Elia that talks about Monolithic vs. Functional GPOs.

> The terms "monolithic" and "functional" refer to how you design them. Monolithic GPOs contain settings from many different areas. For example, a monolithic GPO might contain settings from Administrative Templates, Internet Explorer Maintenance, and Software Installation policies—all within a single GPO. By contrast, functional GPOs typically do one thing. For example, a functional GPO may do only Software Installation or enforce Security settings.

I totally agree with this and my advice to you when trying to decide which to use that your should pick the type of policy configuration that suites your needs.

This also maps very nicely to the 80/20 examples you will see below where you take a more Monolithic approach to the 80% GPO's and more Functional to the 20%. The 80% policies are going to have more setting in them but they will be relatively static where the 20% policies will have fewer settings but probably need to be updated more frequently. This way you should be able to balance the pro's and con's of each policy type in your environment.

REFERENCES

TechNet: Complying with Service Level Agreements

> If you have large or complex GPOs that require frequent changes, consider creating a new GPO that contains only the sections that you update regularly.

Setting (not policies) = Slower SOE

It is often a misconception that splitting up your group policy setting into a lot of Group Policy Objects (GPO's) will slow down Group Policy on your computers. While this might be true if you have many 100's (or thousands) of GPO's applied to your computer this is not normally the reason why computer may slow down processing Group Policies. Normally you will find that its the number of settings you have applied that will cause performance issues and even then you will find that particular setting that will cause more of a performance hit than other. In my experience the policy setting that cause the most likely affect performance are:

1. Printer Mappings (100+)
2. Folder Redirection (Especially with Windows XP and AppData Redirection)

You should also expect that the first time a users logs on with a new account that they should expect a slow logon as the computer will need to apply all policy setting. However subsequent logon's should be much faster as the computer is then only checking the policy is still applied. This is similar to the difference between running a "GPUPDATE" and a "GPUPDATE /FORCE" .

You should also check out the Best Practice for Optimizing Group Policy Performance post by Darren Mar-Elia as this post explains in detail how GPO are applied and what you can do to tweak performance.

While it would be fairly rare to have an environment that has more than a 999 GPO's applied to a single computer still be aware there is a hard limit on the number of GPO's you can apply to any user or computer. Thus trying to keep the number GPO's to a as few as possible is a good idea especially in very large organisations that may uses separate GPO's for installing software packages.

REFERENCE

TechNet: Determining the Number of Group Policy Objects

> Note that a maximum of 999 GPOs is supported for processing GPOs on any one user or computer. If you exceed the maximum, no GPOs will be processed.

Disable User/Computer settings if not in use

If you are creating a GPO that is only meant to be applied to computers (and vice versa for users) then you should disable the unused portion of the GPO. This not only helps guards against accidental change to the section of the GPO that should not be applied it should also give you a small performance boost processing policies on your computers as the GPO does not unnecessarily evaluate parts of the policy that are not configured with any settings.

While I have never seen a performance benefit in disabling the unused portion of a GPO or based on the number of GPO's applied to a computers (see "Settings (not polices) = Slower SOE)" section above) I do encourage that you adhere to these principals to avoid Death of a thousand cuts when it comes to the performance of your systems.

TechNet: Complying with Service Level Agreements

> If a GPO contains only computer or user settings, disable the portion of the policy that does not apply. The destination computer does not scan the portions of a GPO that you disable, which reduces processing time

Avoid using Enforced

In all my time as an Group Policy Administrator I cannot real once a scenario that I required the use of the Enforced feature of Group Policy. At all cost you should avoid this setting as doing so is like using big hammer to a problem that you can probably avoid if designed right.

(RESIST THE URGE)

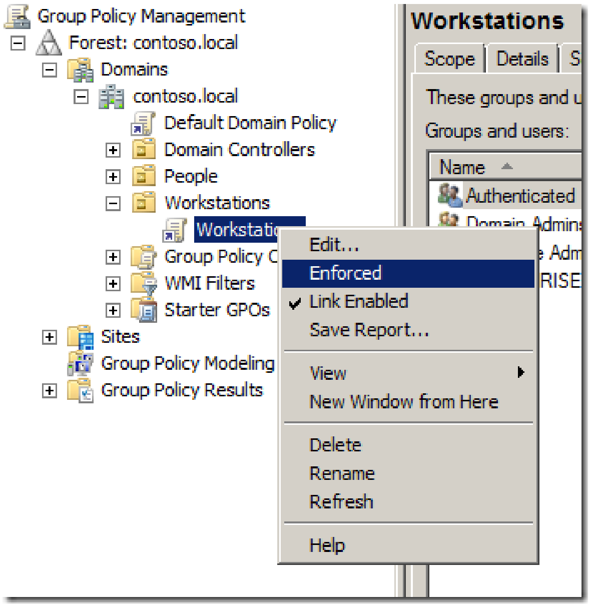

REFERENCES

TechNet: Designing Your Group Policy Model

> Use the **Enforced** and **Block Policy Inheritance** features sparingly. Routine use of these features can make it difficult to troubleshoot policy because it is not immediately clear to administrators of other GPOs why certain settings do or do not apply

Reuse GPO's where possible

If you are in a situation that you want have the same settings you want to apply to all the users or computers in specific OU's your organisation then consider linking the same GPO to these OU's. When naming the GPO chose a name that represents what is common to what you are applying. This is shown in the image below (and in "80/16/4 Example 2") where the policy is named "People Manufacturing" as this is the common two values to where to policy is being applied.

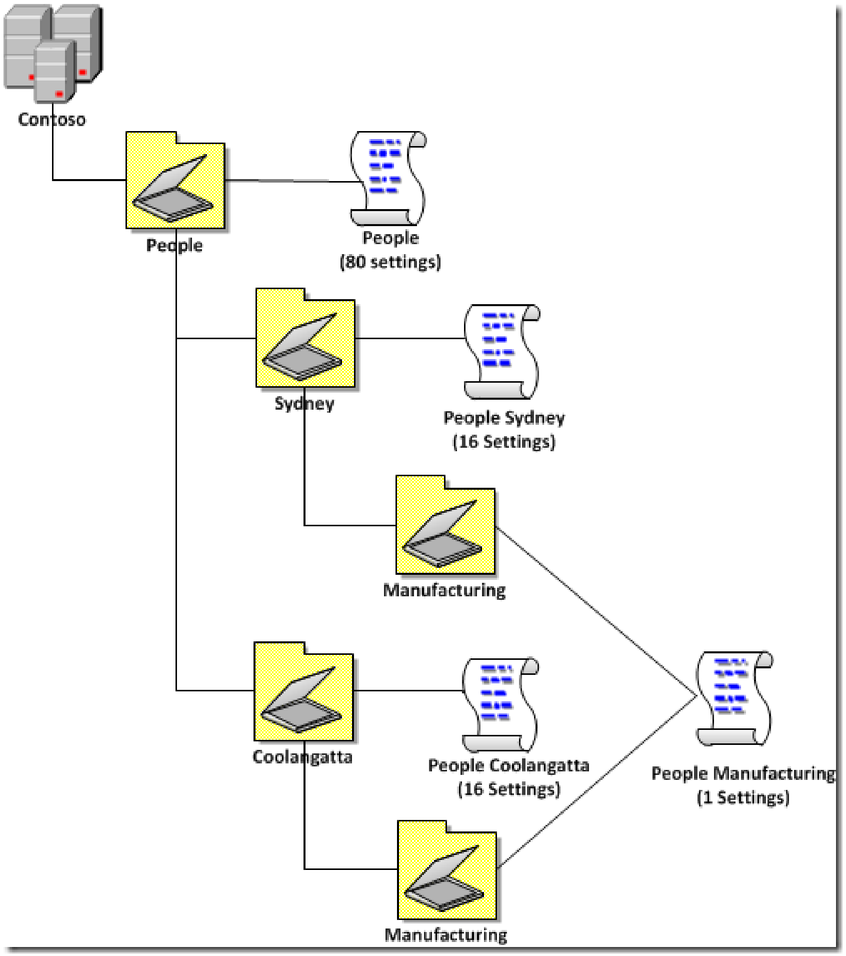

The means the "Sydney" and "Coolangatta" is ignored as that would result in a long policy name of "People Sydney and Coolangatta" Manufacturing". It would be obviously longer again if you had the policy linked to many more sites.

If you have Software Assurance use the Advance Group Policy Management (AGPM) tool

Advanced Group Policy Management (a.k.a. AGPM) is a tool that is available to anyone who is licensed to have Software Assurance. This programs is a change management tool that allows you to check-in and check-out GPO as well as create a list of changes and an audit trail of change to GPO's. You can check out my AGPM install and configuration series at AGPM Part 1: Introduction to Advanced Group Policy Management (a.k.a AGPM) v4. If you have a Group Policy infrastructure of any size or if you have more than one person who is responsible for making changes to GPO's then this is definitely something you should consider.

AGPM is also very good at avoiding GPO editing conflicts as you will find that the "last writer will win" when making policy changes. This means that in an environment that has multiple GPO admins you might find that you could be overwriting each other changes with un-expected results. AGPM gets around this issues as it support the method of checking in and out GPO's for editing meaning that now two GPO administrators can edit a GPO at the same time thus eliminating the possibility of overwriting each other changes.

For even more information on AGPM check out the following links:

Microsoft MDOP Blog
TechNet: Overview of Advanced Group Policy Management
TechNet: A Video tour of Advanced Group Policy Management
TechNet: Technical Overview of AGPM
TechNet: What's New in AGPM

TechNet: Choosing Which Version of AGPM to Install
TechNet: Step-by-Step Guide for Microsoft Advanced Group Policy Management 4.0
TechNet: Operation Guide for Microsoft Advanced Group Policy Management 4.0
Group Policy Blog: Importing and Exporting with AGPM

Create a Test Group Policy Structure

Implement something like AGPM is an excellent way to make sure you have a proper rollback strategy for making changes to Group Policy but sometimes you just want somewhere to test the policy functionality before you put it into production. I would definitely recommend having an isolated replica of the AD structure in for making test however the problem with these environment is that they are normally not a 100% representation of the production environment.

Therefore as a second step in your testing of policy changes before being applied to productions systems you should create a test GP structure that will allow have a selection of users and computers that are in production but are not mission critical. Best to select users that you know are easy to get along with and wont scream to loud when you break something. You can even apply your own computer and users account to this test GP structure but make sure that this is not your only account as you want your computer to still be able to work so you can undo your changes in case you royally stuff something up.

OU METHOD

The image below shows how you could implement a Test OU/GP structure however by creating a separate OU structure to test your group policy. This method provides excellent isolation of your test computers and users to production which may be desired if you want to lessen the impact of any bad configuration changes. However this would mean that you would have the overhead of needing to ensure that all configuration changes to the production GPO's are also replicated to these. Otherwise you may end up with your test environment being configured differently to your production GPO.

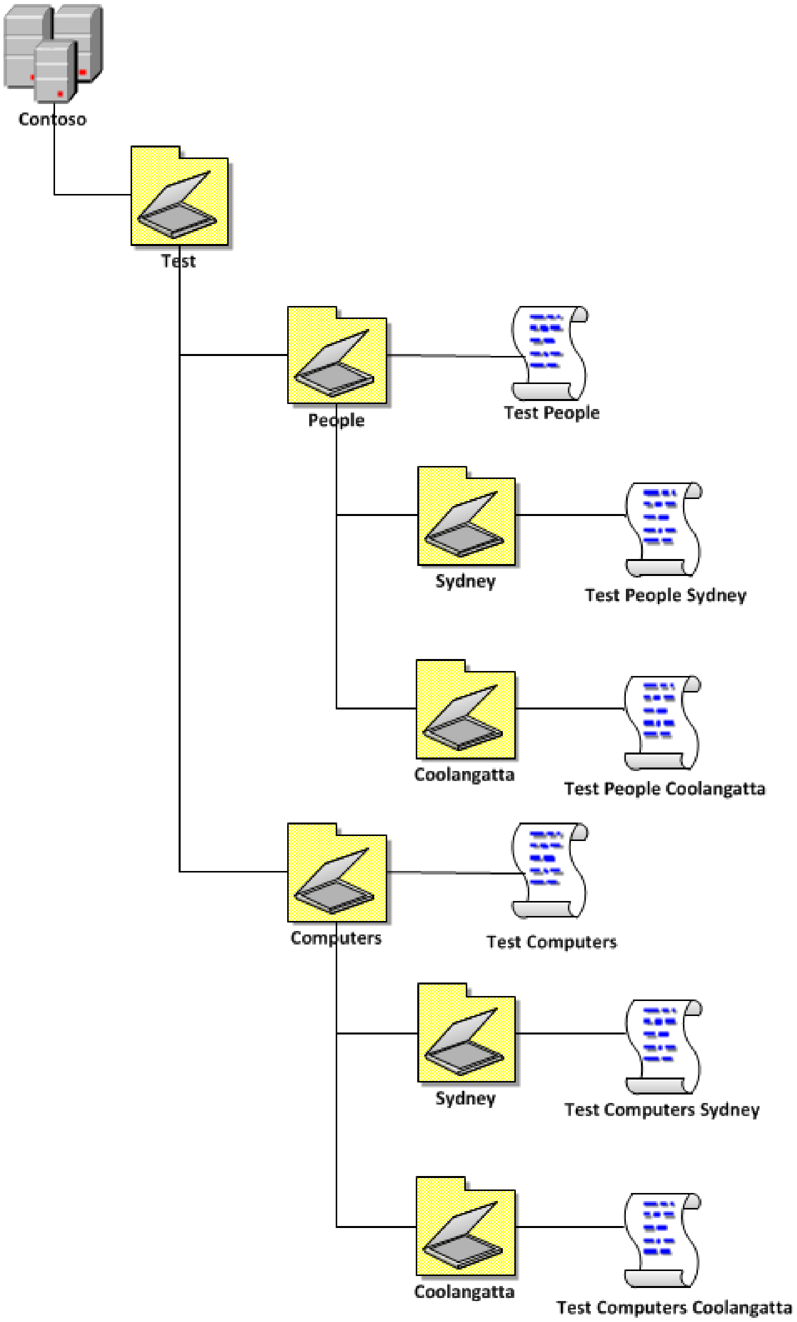

SECURITY GROUP FILTERED METHOD

The Security Group Filtered method applies the test GPO's to the existing OU structure but they are security filtered so they will only apply to the users or computers you want to test. The test GPO will only have the delta configuration changes applied to it for the policy setting that you are testing therefore all other production policies will be implicitly applied to the test objects. Therefore you test computers and users are as close as possible representation of production because they are subject to the production policies. This also mean you do not need to make duplication configuration changes to the GPO's when you do make production changes as the test computers will automatically have the production policies applied. The down side to this method is that unless you are carful in how you apply your security filtering you may inadvertently apply the test changes to your computers users and computers as they are all under the same scope of the test GPO. Another disadvantage of this method is that as you are relying upon security groups to apply the users or computer to the test policy it is possible that you could be a member of multiple test groups and thus be subject to multiple conflicting test GPO's which may make the results somewhat unpredictable.

When not testing GPO changes the Test GPO's should remain configured without any settings and/or the link to the OU should be disable to avoided any extra policy processing overhead to the production users and computers.

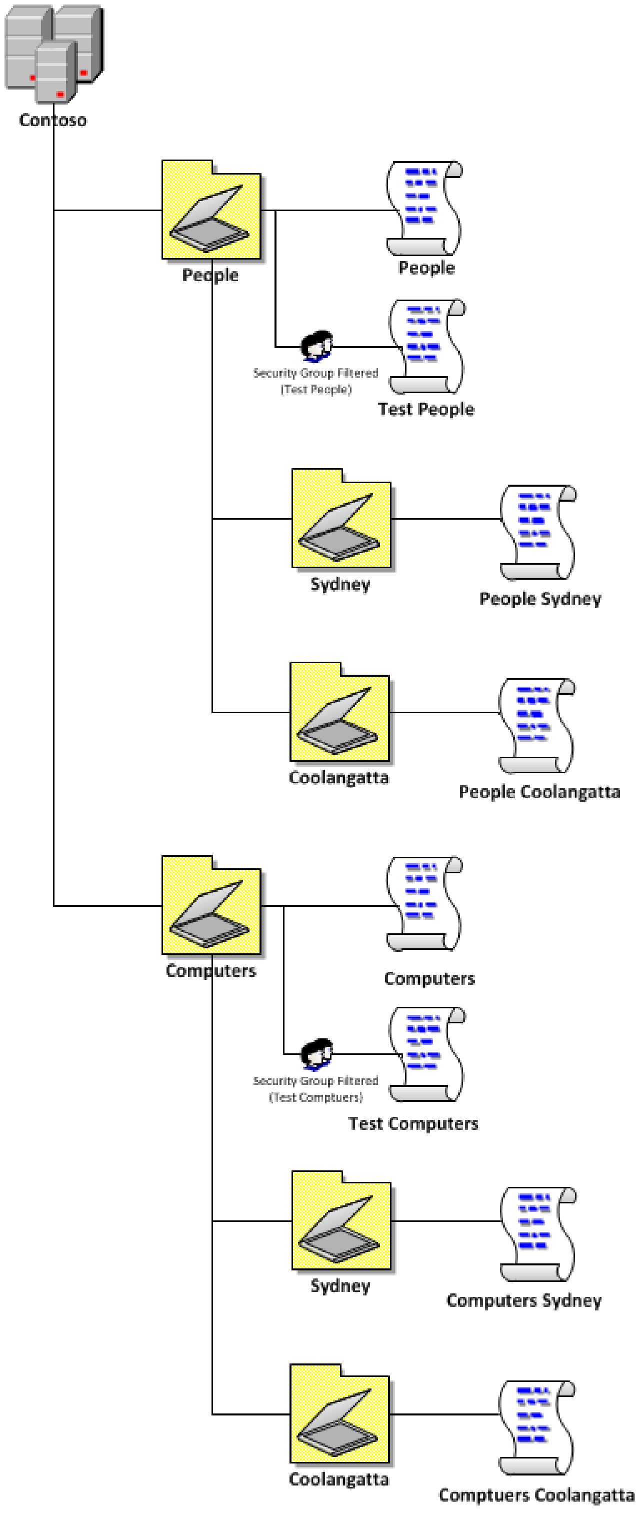

HYBRID METHOD

This method combines both a separate OU structure and separate GPO's but avoids having to use security group filtering. The advantage of this method is that you test environment is still subject to the production GPO's however the test policies are only applied to the users and computers that are located in the Test OU structure. This method totally mitigates accidently applying a test configuration to your production computers and it also eliminates the need to duplicate configuration changes to your production environment.

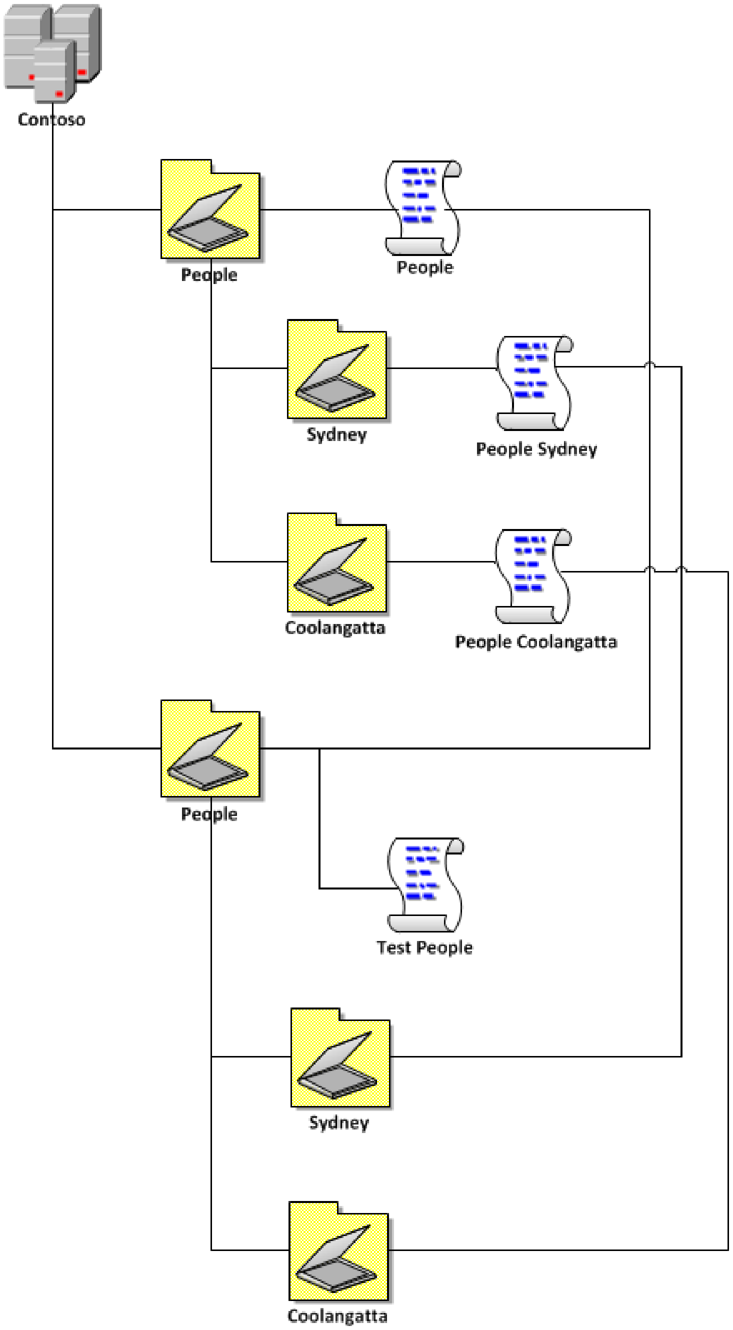

REFERENCE

TechNet: Establishing Group Policy Operational Guidelines

> Always stage Group Policy deployments using the following pre-deployment process

TechNet: Designing Your Group Policy Model

> Prepare a staging environment to test your Group Policy-based management strategy before deploying GPOs into your production environment.

TechNet: Deploying Group Policy

> Always fully test your GPOs in safe (nonproduction) environments prior to production deployment

Backup Often

Especially if you don't have something like AGPM installed in your environment you should seriously consider making a PowerShell script that simple backs up all your new GPO's in your Active Directory every night. Having back up copies of you GPO is very handy especially if you have miss-configured something and you quickly want to rollback to last known good policy setting. For more information on how to do this with PowerShell visit PowerShell Script: Backup all GPOs that have been modified this month from the Group Policy Team Blog.

REFERENCES

TechNet: Defining Group Policy Operational Procedures

> You should also create regular backups of your GPOs

Edit Default Domain Policies Sparingly

Unless you are changing the default domain password policy then it is strongly recommended that you do not modify the Default Domain or Default Domain Controller group policy objects as making a mistake in these two policies up can really mess up your Active Directory. If you want to make a change to all your DC or your entire domain then consider making a separate new group policy at the same level as the default policies. This will at least allow you to un-do any change selectively disabling the offending policies if something does go wrong.

REFERENCE

TechNet: Linking GPOs

> If you need to modify some of the settings contained in the **Default Domain Policy GPO**, it is recommended that you create a new GPO for this purpose, link it to the domain, and set the **Enforce** option. In general, do not modify this or the **Default Domain Controller Policy GPO**. If you do, be sure to back up these and any other GPOs in your network by using GPMC to ensure you can restore them.

TechNet: Establishing Group Policy Operational Guidelines

> Do not modify the default domain policy or default domain controller policy unless necessary. Instead, create a new GPO at the domain level and set it to override the default settings in the default policies.

Update: Here is another post I have found that confirms this http://jorgequestforknowledge.wordpress.com/2011/10/23/best-practices-for-the-default-domain-policy-and-the-default-domain-controllers-policy-gpos/

Avoid using Group Policy Software Assignment

I know it sounds strange for a Group Policy expert to say avoid using Group Policy but this is definite one case where you should consider using other software deployment products due to their vastly superior features.

Group Policy Software Installations (a.k.a. GPSI) is a way you can deploy an MSI based application to your computers using Group Policy. This can be very useful way of deploying a standard set of applications to your computers however when compared to the advanced targeting features of SCCM software deployment or App-V this limitations of this method of software deployment quickly becomes evident.

One common problem I see when deploying software this way is the "Un-install when falls out of scope" options. This can be very handy when you want to move a computer to another OU and you want all the software packages that are not needed any more to un-install. This is even worse when you try to move an computer between domains as the computer will then un-install and re-install all the applications assigned to it which can take a VERY LONG time. Even when you have the "Un-install when falls out of scope" not ticked on the source domain and you move the computer to a new domain you will find that the installer service will still need to do a repair/check install of all the applications of the new domain even if the applications are already installed. However this also means that when the computer is removed from a domain then you have to wait for all the application's to un-install during the next reboot. The un-installing of application can obviously take a long time if you have many applications install via this method. If you don't select this options then you will find that your computer will over time build up the a number of installed applications installed on your computers that will affect performance, stability and licensing costs. The other inflexibility of doing software assignment to the computers via GPSI is that they will only install on the next reboot of the computer. Meaning that a user will need to do a full reboot of their computer before they will be able to start using the new applications.

The other restriction of GPSI is that you are limited to deploying only Microsoft Software Install (a.k.a. MSI) packages. Where tools like SCCM and App-V will allow you to deploy application via a silent command line option or via a sequenced application.

So due to all these targeting issues with GPSI software then I strongly recommend that you consider using either Microsoft SCCM package deployment or Microsoft App-V due to the superior targeting and features these products offer. For more information on the advantages of Microsoft App-V then i strongly recommend that you checkout the series of App-V FAQ at http://blog.stealthpuppy.com/tag/appvfaq .

REFERENCES

Office 2007 Deployment via Group Policy

> Office 2007 is no longer deployed using transform files

Below are the only scenarios that should be used when deploying Office 2007 via GPSI. While this article is specific to Office 2007 I would also say that the same limitations should be used when considering GPSI for other applications as well.

TechNet: Use Group Policy Software Installation to deploy the 2007 Office system

> You can use the Software Installation extension of Group Policy to deploy the 2007 Office system to *computers* if the following conditions exist:
>
> - Small organizations that have already deployed and configured Active Directory
> - Organizations or departments that comprise a single geographic area
> - Organizations with consistent hardware and software configurations on both clients and servers

Never edit Group Policy Objects from the Domain Controller

To often I see people editing their GPO's directly from a Domain Controller in their organisation as they are not aware they can do this remotely. The Remote Server Admins Tools (a.k.a. RSAT) have will give you the option to install (See instructions here) the Group Policy Management Console on any workstation or server running Vista/2008 or greater. I strongly encourage you to do this as if you are performance day to day management of your active directory (e.g. Creating users, editing Group Policy and adding/removing users from groups) then sooner or later you will find that you might affect the stability of your DC (which would be BAD).

Apply policies as high as possible

When given the choice of applying the same policy at multiple lower locations or just one locations higher always try to link the policy as high up as possible in the OU tree. If there are cases where you want to apply the policy setting at all levels except for a minority of the lower sub-OU's then simple apply a different policy on the fewer OU's to make the exception.

Bad Example	**Good Example**
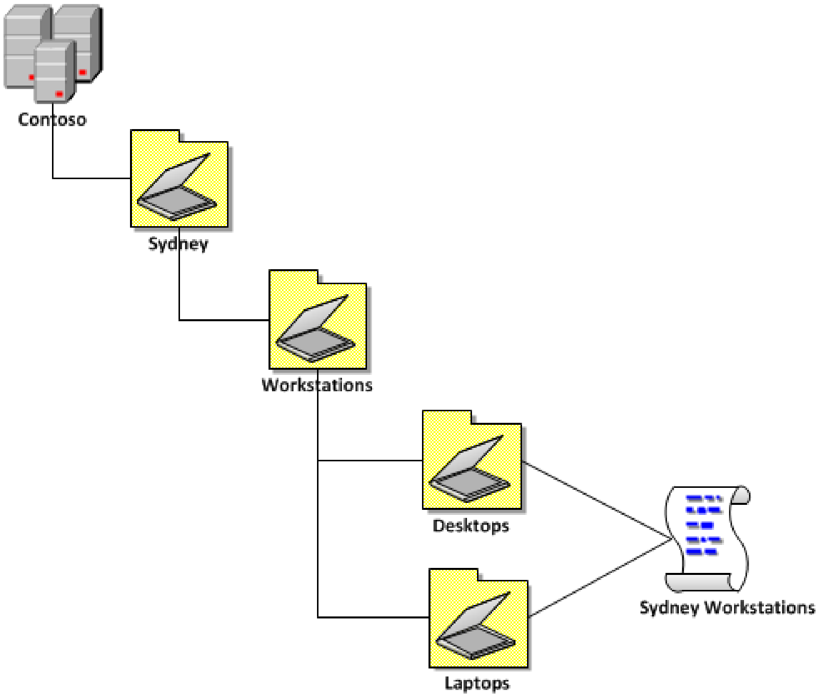 Contoso Sydney Workstations Desktops Laptops Sydney Workstations	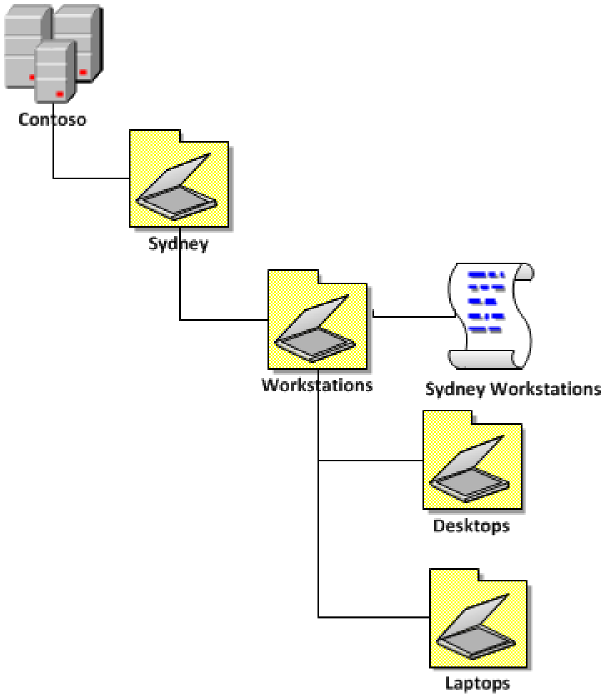 Contoso Sydney Workstations Sydney Workstations Desktops Laptops