DOWNLOAD PDF OPTIMIZING ACTIVE DIRECTORY 26.

Chapter 1 : News, Tips, and Advice for Technology Professionals - TechRepublic

Right-click on Active Directory Domain Services and click Start The defragmentation of the Active Directory database is now complete. Tags Active Directory active directory tutorial How-To Migration MVP Step-By-Step tutorial Windows Server Windows Server Windows Server Windows Server R2 Windows Server Windows Server R2.

Whenever you make an update to Active Directory, your change is added to the Active Directory database on a domain controller, which then replicates the change to all of the other domain controllers in the domain. The more domain controllers that are in the domain, the more replication-related traffic you can expect on your network. Fortunately, there are some things you can do to get a handle on replication-related traffic. This requires you to divide the domain into sites and then schedule replication between the sites. The problem is that not everyone has his or her network subnetted. The only way to get the performance gains associated with site segregation is to ensure that your network is effectively divided into subnets. Sites do two main things for Windows. First, they allow you to schedule replication rather than having it occur on an as-needed basis. Second, they condense replication traffic into single sets. To see why scheduling replication-related traffic is important, consider this: If someone at a satellite office were to make an update to Active Directory, the update would most likely be related to an Active Directory object that applies directly to that office, such as a password change for a user in that office. If a user were to change his or her password, it would be necessary for the password change to eventually be replicated to every domain controller in the domain. Sites make this selective updating possible. For example, suppose you established a site replication schedule of 15 minutes. But it could be up to 15 minutes before domain controllers in remote sites were updated. If the remote office had five domain controllers, then five sets of replication traffic would flow to the remote office for every Active Directory update. When sites are used, a single set of replication traffic flows between the two sites. You must have at least one site connector connecting any two sites, although it is possible to use redundant or transitive site connectors. Although there are a lot of configuration options related to site connectors, the two most important are the cost and the replication frequency. The cost is simply a numeric value that Windows uses to determine which site connector to use. If a site had two site connectors to another site, and one had a cost of 1 and the other had a cost of 2, the connector with the lowest cost would always be used. The higher cost connection would be used only if the lower cost connection was unavailable. The replication frequency parameter controls the amount of time between replication updates. The minimum value is 15 minutes, and the maximum is 10, minutes, or one week. The main point to remember is that longer replication frequencies mean better performance but fewer updates. Lower replication frequencies mean a more consistent Active Directory database, but they also mean you will be sacrificing some bandwidth to get those frequent updates. On smaller networks, minute replication cycles are acceptable. On larger networks, use minute replication cycles. Of course, these are just guidelines.

DOWNLOAD PDF OPTIMIZING ACTIVE DIRECTORY 26.

Chapter 2 : Optimize Active Directory for Azure AD, Office and the Cloud | Gartner Webinars

Active Directory Database Optimization February 19, by Dishan M. Francis 1 Comment Like any other database active directory database also get fragmented as its write and retrieve data from the database.

February 19, by Dishan M. Francis 1 Comment Like any other database active directory database also get fragmented as its write and retrieve data from the database. It will also grow on size without clearing unused hard drive space. It needs to have regular optimization of active directory database to have better performances. How we can do it? In windows OS we uses the defragment tool to optimize the computer hard drive. There is similar procedure we can use to defrag active directory database. There are two type of defragmentation use with active directory database. Online Defragmentation With windows serer Microsoft introduced this method. It is runs in certain intervals default is every 12 hours automatically to defrag active directory database. It is part of active directory garbage collection process. It will optimize the data storage and reclaims the space for new active directory objects. But this will not reduce the size of the active directory database. The important thing is it not required to bring any service offline to do this. Offline Defragmentation As the name says to do this process we need stop the active directory service. To do this system will create compact version of the existing active directory database in different location. Once process is created the new defragmented database it will copy the compact version in to the original location. This is the same tool we can use to check for the active directory errors. Since Ad service will go down you need to measure how it will affect company operations. The time it will take depends on the size of the AD database and the how bad it fragmented. Click yes to continue. For demo I created folder C: The time it will take depends on the size of the database. To complete the process as screen says copy the defragmented database from C: After that we have successfully defrag the AD database. Now go to Services. If you have any question regarding the article feel free to contact me on rebeladm live.

Chapter 3 : Optimizing active directory population in C# (blog.quintoapp.comoryServices)

For more info on how Active Directory Searches work, see How Active Directory Searches Work. Some scenarios in which to add indices include: Client load in requesting the data is generating significant CPU usage and the client query behavior cannot be changed or optimized.

When an attribute is indexed, the index applies to each object to which the attribute is associated. You cannot place an index on an attribute and have it apply to only one class of object. To specify the index for an attribute, you must modify the searchFlags property for the attribute in the schema. The searchFlags Property The searchFlags property is part of the definition for a schema attribute and is composed of bit values. These bit values determine how an attribute is handled by Active Directory. The first three bits control how an attribute is indexed. By setting the first bit of the property to one, a database-wide index is created for the attribute. Setting the second bit to one creates an index on each container that holds an object that uses the attribute. This kind of index helps one-level searches, and Windows XP defines the supported sort orders for one-level VLV searches. If the third bit is set, then the attribute is included in the filter expansion of ANR searches of the ANR attribute set only if the attribute is already normally indexed by setting the first bit to one. In the Windows Server family, if the sixth bit is set, then a substring index is created on the attribute, allowing efficient medial substring queries to be performed. Note that a substring index is significantly more expensive to create and maintain, so update speeds will be affected. Ambiguous name resolution Ambiguous name resolution ANR helps find a user object when a unique identifying value is not known. Although ANR was designed for use with user objects, it can be used on any type of object. Therefore, you may need to indicate the class in the search as well. Determining when the index has changed When the searchFlags property is modified, a background task starts to create the index. This task may take a while to complete depending on which attribute is indexed and how many objects are using the attribute. To be certain that the index has been created, look at the event log for the following message: When one of the bit values is set to zero to turn off the index, a background task is started for removing the index. To be certain that the index has been removed, look at the event log for the following message: The statistics control reports the following information: Number of actual database operations generated by the LDAP search operation. Number of entries evaluated by the optimizer. Time needed by Directory Services to perform the operation. Indexes used by the optimizer for filter evaluation. This control is available to any application that uses LDAP. On the Options menu, click Controls. In the Object Identifier box, enter 1.

Chapter 4 : Optimize your Active Directory environment with Azure Log Analytics | Microsoft Docs

One of the biggest drawbacks to Active Directory is its distributed nature. Whenever you make an update to Active Directory, your change is added to the Active Directory database on a domain.

Must be online when schema updates are performed. Domain Naming Master Used to add and to remove domains and application partitions to and from the forest. Must be online when domains and application partitions in a forest are added or removed. Primary Domain Controller Domain Receives password updates when passwords are changed for the computer and for user accounts that are on replica domain controllers. Consulted by replica domain controllers that service authentication requests that have mismatched passwords. Default target domain controller for Group Policy updates. Target domain controller for legacy applications that perform writable operations and for some admin tools. Must be online and accessible 24 hours a day, seven days a week. Must be online for newly promoted domain controllers to obtain a local RID pool that is required to advertise or when existing domain controllers have to update their current or standby RID pool allocation. Infrastructure Master Domain Application partition Updates cross-domain references and phantoms from the global catalog. For more information, click the following article number to view the article in the Microsoft Knowledge Base: This placement is frequently correct for directories that have just a few domain controllers. In a directory that has many domain controllers, the default placement may not be the best match for your network. Consider the following in your selection criteria: Place roles on domain controllers that are can be accessed by the computers that need access to a given role, especially on networks that are not fully routed. If a role has to be moved to a different domain controller, and the current role holder is online and available, you should transfer not seize the role to the new domain controller. FSMO roles should only be seized if the current role holder is not available. For more information, go to the following Microsoft website: If the role holder can be made operational before the role is needed, you may delay seizing the role. If role availability is critical, transfer or seize the role as required. The PDC role in each domain should online at all times. Select a direct intrasite replication partner for existing role holders to act as a standby role holder. If the primary owner goes offline or fails, transfer or seize the role to the designated standby FSMO domain controller as required. Place the domain naming master on the forest root PDC. The addition or removal of domains should be a tightly controlled operation. Place this role on the forest root PDC. Certain operations that use the domain naming master, such as creating or removing domains and application partitions, fail if the domain naming master is not available. On a domain controller that runs Microsoft Windows , the domain naming master must also be hosted on a global catalog server. On domain controllers that run Windows Server or later versions, the domain naming master does not have to be a global catalog server. Place the PDC on your best hardware in a reliable hub site that contains replica domain controllers in the same Active Directory site and domain. In large or busy environments, the PDC frequently has the highest CPU utilization because it handles pass-thru authentication and password updates. If high CPU utilization becomes a problem, identify the source, and this includes applications or computers that may be performing too many operations transitively targeting the PDC. Techniques to reduce CPU include the following: All domain controllers in a particular domain, and computers that run applications and admin tools that target the PDC, must have network connectivity to the domain PDC. RID master overhead is light, especially in mature domains that have already created the bulk of their users, computers, and groups. The domain PDC typically receives the most attention from administrators. Therefore, co-locating this role on the PDC helps ensure reliable availability. Make sure that existing domain controllers and newly promoted domain controllers, especially those promoted in remote or staging sites, have network connectivity to obtain active and standby RID pools from the RID master. Legacy guidance suggests placing the infrastructure master on a non-global catalog server. There are two rules to consider: In a forest that contains a single Active Directory domain, there are no phantoms. Therefore, the infrastructure master has no work to do. The infrastructure master may be placed on any domain controller in the domain, regardless of whether that domain controller hosts the global catalog or not. If every domain controller in a domain that is part of a multidomain forest also hosts the global catalog,

there are no phantoms or work for the infrastructure master to do. The infrastructure master may be put on any domain controller in that domain. In practical terms, most administrators host the global catalog on every domain controller in the forest. If every domain controller in a given domain that is located in a multidomain forest does not host the global catalog, the infrastructure master must be placed on a domain controller that does not host the global catalog.

Chapter 5 : Optimizing authentication between NPS and Active directory server

Tips for optimizing your Active Directory before your move How Quest can help you prepare for, migrate to, secure and manage your new environment See for yourself how to make your move to Windows Server simpler than you imagined so you can begin seeing a return on your investment.

Understanding how recommendations are prioritized Every recommendation made is given a weighting value that identifies the relative importance of the recommendation. Only the 10 most important recommendations are shown. How weights are calculated Weightings are aggregate values based on three key factors: The probability that an issue identified causes problems. A higher probability equates to a larger overall score for the recommendation. The impact of the issue on your organization if it does cause a problem. A higher impact equates to a larger overall score for the recommendation. The effort required to implement the recommendation. A higher effort equates to a smaller overall score for the recommendation. The weighting for each recommendation is expressed as a percentage of the total score available for each focus area. Focus areas Security and Compliance - This focus area shows recommendations for potential security threats and breaches, corporate policies, and technical, legal and regulatory compliance requirements. Availability and Business Continuity - This focus area shows recommendations for service availability, resiliency of your infrastructure, and business protection. Upgrade, Migration and Deployment - This focus area shows recommendations to help you upgrade, migrate, and deploy Active Directory to your existing infrastructure. The recommendations are based on the knowledge and experiences gained by Microsoft engineers across thousands of customer visits. However, no two server infrastructures are the same, and specific recommendations may be more or less relevant to you. For example, some security recommendations might be less relevant if your virtual machines are not exposed to the Internet. Some availability recommendations may be less relevant for services that provide low priority ad hoc data collection and reporting. Issues that are important to a mature business may be less important to a start-up. You may want to identify which focus areas are your priorities and then look at how your scores change over time. Every recommendation includes guidance about why it is important. You should use this guidance to evaluate whether implementing the recommendation is appropriate for you, given the nature of your IT services and the business needs of your organization. Use health check focus area recommendations After it is installed, you can view the summary of recommendations by using the Health Check tile on the solution page in the Azure portal. View the summarized compliance assessments for your infrastructure and then drill-into recommendations. To view recommendations for a focus area and take corrective action Click the Overview tile for your Log Analytics workspace in the Azure portal. On the Health Check page, review the summary information in one of the focus area blades and then click one to view recommendations for that focus area. On any of the focus area pages, you can view the prioritized recommendations made for your environment. Click a recommendation under Affected Objects to view details about why the recommendation is made. You can take corrective actions suggested in Suggested Actions. When the item has been addressed, later assessments records that recommended actions were taken and your compliance score will increase. Corrected items appear as Passed Objects. Ignore recommendations If you have recommendations that you want to ignore, you can create a text file that Log Analytics will use to prevent recommendations from appearing in your assessment results. To identify recommendations that you will ignore In the Azure portal on the Log Analytics workspace page for your selected workspace, click the Log Search tile. Use the following query to list recommendations that have failed for computers in your environment. Choose recommendations that you want to ignore. To create and use an IgnoreRecommendations. Paste or type each RecommendationId for each recommendation that you want Log Analytics to ignore on a separate line and then save and close the file. Put the file in the following folder on each computer where you want Log Analytics to ignore recommendations. You can use the following Log Search queries to list all the ignored recommendations. The check runs every seven days. Is there a way to configure how often the health check runs? Not at this time. If a server is decommissioned, when will it be removed from the health check? If a server does not submit data for 3 weeks, it is removed. What is the name

of the process that does the data collection? The actual data collection on the server takes about 1 hour. It may take longer on servers that have a large number of Active Directory servers. Is there a way to configure when data is collected? Why display only the top 10 recommendations? Instead of giving you an exhaustive overwhelming list of tasks, we recommend that you focus on addressing the prioritized recommendations first. After you address them, additional recommendations will become available. If you prefer to see the detailed list, you can view all recommendations using Log Search. Is there a way to ignore a recommendation? Yes, see Ignore recommendations section above.

Chapter 6 : Active Directory in Exchange organizations | Microsoft Docs

Integration with the cloud compels many organizations to re-evaluate their IAM approach to Active Directory with the goal of achieving greater simplicity and consistency. This assessment guides identity architects through common approaches to AD optimization in enterprise-only and hybrid scenarios.

Troubleshooting Capacity planning Properly deploying a sufficient number of domain controllers, in the right domain, in the right locales, and to accommodate redundancy is critical to ensuring servicing client requests in a timely fashion. This is an in-depth topic and outside of the scope of this guide. Updates and evolving recommendations Massive improvements in both Active Directory and client performance optimizations have occurred over the last several generations of the operating system and these efforts continue. We recommend that the most current versions of the platform be deployed to get the benefits, including: We recommend keeping current with these updates. Hardware basics This is a summary of key points covered in much greater depth in the Capacity Planning for Active Directory Domain Services and is not a replacement for that content. Read the following sections to optimize hardware for responsiveness of domain controllers to client requests. Avoid going to disk Active Directory caches as much of the database as memory allows. Fetching pages from memory are orders of magnitude faster than going to physical media, whether the media is spindle or SSD based. For more info about storage subsystem tuning, see Performance Tuning for Storage Subsystems. Active Directory Best Practices recommend putting enough RAM to load the entire DIT into memory, plus accommodate the operating system and other installed applications, such as anti-virus, backup software, monitoring, and so on. For limitations of the legacy platforms, see Memory usage by the Lsass. Put the operating system, logs, and the database on separate volumes. If all or the majority of the DIT can be cached, once the cache is warmed and under a steady state, this becomes less relevant and offers a little more flexibility in storage layout. In scenarios where the entire DIT cannot be cached, the importance of splitting the operating system, logs, and database on separate volumes becomes more important. Write-heavy scenarios do not greatly benefit from the Active Directory cache. To guarantee the transactional durability of data that is written to the directory, Active Directory does not perform disk write caching. Instead, it commits all write operations to the disk before it returns a successful completion status for an operation, unless there is an explicit request not to do this. The following are hardware recommendations that might improve performance for these scenarios: Most Active Directory scenarios are predominantly read-based, thus the statistics on the volume hosting the DIT are the most important to inspect. However, do not overlook monitoring the rest of the drives, including the operating system, and log files drives. To determine if the domain controller is properly configured to avoid storage being the bottleneck for performance, reference the section on Storage Subsystems for standards storage recommendations. Across many environments, the philosophy is to ensure that there is enough head room to accommodate surges or spikes in load. These thresholds are warning thresholds where the head room to accommodate surges or spikes in load becomes constrained and client responsiveness degrades. In short, exceeding these thresholds is not bad in the short term 5 to 15 minutes a few times a day , however a system running sustained with these sorts of statistics is not fully caching the database and may be over taxed and should be investigated. To maintain consistency of data, all changes must be written to the log. There is no good or bad number here, it is only a measure of how much the storage is supporting. Across many environments, the philosophy is to ensure that there is enough head room to accommodate surges or spikes in load to minimize impact on client responsiveness in these scenarios. Systems spending sustained periods above the thresholds should be investigated to how to reduce processor loads. For more info on how to select a processor, see Performance Tuning for Server Hardware. Add hardware, optimize load, direct clients elsewhere, or remove load from the environment to reduce CPU load. Avoid overloading the network adapter Just like with processors, excessive network adapter utilization will cause long wait times for the outbound traffic to get on to the network. Active Directory tends to have small inbound requests and relatively to significantly larger amounts of data returned to the client systems. Sent data far exceeds received data. This threshold is a warning threshold where the head room to accommodate surges or spikes in load becomes

constrained and client responsiveness degrades. In short, exceeding these thresholds is not bad in the short term 5 to 15 minutes a few times a day , however a system running sustained with these sorts of statistics is over taxed and should be investigated. For more info on how to tune the network subsystem, see Performance Tuning for Network Subsystems. Proper placement of domain controllers and site considerations Proper site definition This is critical to performance. Clients falling out of site can experience poor performance for authentications and queries. The operating system prefers IPv6 to IPv4 when both are configured. This can cause exhaustion of the ATQ Thread Pool and cause the domain controller to become unresponsive. The appropriate resolution to this is to properly define the site topology for IPv6. As a workaround, you can optimize the name resolution infrastructure to respond quickly to domain controller requests. Optimize for referrals Referrals are how LDAP queries are redirected when the domain controller does not host a copy of the partition queried. When a referral is returned, it contains the distinguished name of the partition, a DNS name, and a port number. The client uses this information to continue the query on a server that hosts the partition. This is a DCLocator scenario and all of the recommendations site definitions and domain controller placement is maintained, but applications which depend on referrals are often overlooked. It is recommended to ensure AD Topology including site definitions and domain controller placement properly reflects the needs of the client. Also, this may include having domain controllers from multiple domains in a single site, tuning DNS settings, or relocating the site of an application. Optimization considerations for trusts In an intra-forest scenario, trusts are processed according to the following domain hierarchy: This means that secure channels at the forest root, and each parent, can become overloaded due to aggregation of authentication requests transiting the DCs in the trust hierarchy. This may also incur delays in Active Directories of large geographical dispersion when authentication also has to transit highly latent links to affect the above flow. Overloads can occur in inter-forest and down-level trust scenarios. The following recommendations apply to all scenarios: Create shortcut trusts as appropriate based on load. Ensure that every domain controller in the domain is able to perform name resolution and communicate with the domain controllers in the trusted domain. Ensure locality considerations are taken into account for trusts. Enable Kerberos where possible and minimize use of the secure channel to reduce risk of running into MaxConcurrentAPI bottlenecks. Cross domain trust scenarios are an area that has been consistently a pain point for many customers. Name resolution and connectivity issues, often due to firewalls, cause resource exhaustion on the trusting domain controller and impact all clients. Furthermore, an often overlooked scenario is optimizing access to trusted domain controllers. The key areas to ensure this works properly are as follows: Ensure the DNS and WINS name resolution that the trusting domain controllers are using can resolve an accurate list of domain controllers for the trusted domain. Statically added records have a tendency to become stale and reintroduce connectivity problems over time. Ensure proper configuration of forwarders, conditional forwards, and secondary copies for both forward and reverse lookup zones for every resource in the environment which a client may need to access. Again, this requires manual maintenance and has a tendency to become stale. Consolidation of infrastructures is ideal. Domain controllers in the trusting domain will attempt to locate domain controllers in the trusted domain that are in the same site first and then failback to the generic locators. Converge site names between the trusted and trusting domains to reflect domain controller in the same location. Ensure subnet and IP address mappings are properly linked to sites in both forests. Ensure ports are open, according to DCLocator needs, for domain controller location. If firewalls exist between the domains, ensure that the firewalls are properly configured for ALL trusts. If firewalls are not open, the trusting domain controller will still attempt to access the trusted domain. If communication fails for any reason, the trusting domain controller will eventually time out the request to the trusted domain controller. However, these time outs can take several seconds per request and can exhaust network ports on the trusting domain controller if the volume of incoming requests is high. The client may experience the waits to timeout at the domain controller as hung threads, which could translate to hung applications if the application runs the request in the foreground thread. For more info, see How to configure a firewall for domains and trusts. Use DnsAvoidRegisterRecords to eliminate poorly performing or high-latency domain controllers, such as those in satellite sites, from advertising to the generic locators. Note There is a practical limit of about 50 to the number of domain controllers the client can

consume. These should be the most site-optimal and highest capacity domain controllers. Consider placing domain controllers from trusted and trusting domains in the same physical location. For all trust scenarios, credentials are routed according to the domain specified in the authentication requests. When the domain parameters for these APIs are passed a NULL value, the domain controller will attempt to find the account name specified in every trusted domain available.

Chapter 7 : Optimizing DFS Referrals: SiteCostedReferrals and PreferLogonDC â€“ AD Troubleshooting

This video shows how to perform offline defragmentation of Active Directory Database in Windows Server R2.

Chapter 8 : OPTIMIZE ACTIVE DIRECTORY REPLICATION

An Active Directory Optimization Reference Architecture (ADORA) defining the objective end-state in terms of principles, rules, patterns, and technical positions. Support for the delivery of near-term AD optimization capability enhancements within the.

Chapter 9 : Performance Tuning for Active Directory Servers - Windows 10 hardware dev

The Active Directory (AD) database is constantly changing, and over time, these changes can cause the database to respond to the system more slowly than necessary.