



# Multi-Tenant Scalability Guidance for Exchange Server 2010 Service Pack 2

---

Customer Advisory Team  
Exchange Product Group  
Microsoft Corporation

December 2011

**Contents**

- Introduction ..... 3
- Scalability Guidance ..... 3
  - Tenant Scalability ..... 3
  - Transport Rules ..... 4
  - Accepted Domains ..... 4
  - Public Folders and OAB Distribution ..... 5
  - General Recommendations for High-Scale Deployments ..... 5
  - Active Directory Optimizations ..... 6
- Exchange Server 2010 SP2 Multi-Tenant Scalability Testing ..... 6
  - Goals ..... 6
  - Testing Methodology ..... 7
  - Scale Testing Lab Setup ..... 8
    - Distribution of Virtual Server Images on Hyper-V Root Servers ..... 8
    - Configuration of Virtual Server Images (RAM, CPU, Disk) ..... 9
  - Detailed Scale Testing Results ..... 10
    - Test Results at Measurement Points ..... 10
    - Results at Scale Target ..... 10
    - Validation with Simulated Client Load ..... 10
  - Active Directory Sizing Observations ..... 11
- Summary ..... 12
- Legal Notice ..... 12

## Introduction

Service Pack 2 (SP2) for Microsoft Exchange Server 2010 introduced the Address Book Policies (ABP) feature which allows segmentation of directory information. Use of the ABP feature along with other supported configuration changes provides a platform for building a custom multi-tenant Exchange 2010 SP2 deployment. Various caveats apply when considering use of ABPs in a multi-tenant or “hosting” environment. In [Multi-Tenancy and Hosting Guidance for Exchange Server 2010 SP2](#), Microsoft describes the proper implementation of a multi-tenant solution using Exchange Server 2010 SP2 and describes the various caveats associated with such an implementation.

This document describes Microsoft recommendations for scaling a multi-tenant deployment of Exchange Server 2010 SP2 using the methods described in the Multi-Tenancy and Hosting Guidance for Exchange Server 2010 SP2 white paper. This document is not applicable to deployments that use the /hosting installation switch. That switch installs a modified version of Exchange that uses a different directory schema for tenant isolation. Microsoft has announced that the /hosting option will not be carried forward into the next release of Exchange, and while it continues to be supported in Exchange Server 2010 SP1 and SP2, no further updates or improvements will be made. Therefore, Microsoft is providing deployment guidance for building multi-tenant deployments using Exchange 2010 SP2.

Given past history with Exchange 2007-based hosting platforms such as the Microsoft Hosted Messaging and Collaboration (HMC) platform as well as an understanding of Exchange 2010 architectural changes, this document provides guidance on the scalability of specific areas of concern with Exchange 2010 SP2. This document does not attempt to cover all potential scalability limits in the product, and any high-scale deployment should proceed with necessary caution during the planning, testing, deployment, and scaling phases. The testing that was performed in Microsoft labs is described here to provide context to the provided recommendations.

## Scalability Guidance

### Tenant Scalability

In Exchange Server 2010 SP2, the definition of a “tenant” is somewhat loose. There is no Exchange cmdlet which can be run that results in the creation of a tenant. For the purposes of this document, we will assume that a tenant is a hosted organization configured with a set of administrative policy objects, such as an address book policy (ABP), along with a set of users.

Scale testing performed in Microsoft labs has demonstrated that the Exchange 2010 SP2 build is capable of functional operation with 50,000 tenants configured in a single Active Directory forest. Microsoft recommends not exceeding 50,000 tenants in a single forest. Customers who need to deploy capacity for greater than 50,000 tenants should design a multi-forest architecture with a provisioning system that is capable of supporting multiple Exchange capacity forests.

The tenants used in this testing included the following (additional details of this configuration are available in the testing section later in this document):

- A tenant organizational unit in Active Directory

- Two global security groups

*Note: Global security groups were used for the purpose of applying permissions to Active Directory objects outside of the Exchange system. As a best practice, universal groups should be used with Exchange 2010 as they can be managed using the Exchange management tools.*

- An accepted SMTP domain
- A global address list (GAL)
- An “All Rooms” address list
- An “All Users” address list
- An “All Contacts” address list
- An “All Groups” address list
- An offline address book (OAB)
- An address book policy (ABP)
- Three mailbox-enabled users
- Two mail-enabled contacts

Many multi-tenant messaging deployments contain a large number of tenants and a small number of users per tenant. This test was designed to evaluate the scalability of an Exchange 2010 SP2 solution in this type of configuration, specifically validating the overall number of tenants that could be successfully provisioned in a single Exchange organization without focusing on the overall number of users configured per-tenant or globally within the system. Administrators interested in deployments which contain a large number of users per tenant can also benefit from this guidance. However, they should pay careful attention to the overall impact of the total number of users on the scalability of the system as that may become a more limiting factor than the total number of tenants.

## Transport Rules

Based on testing performed for the creation of this document, as well as sizing and scalability of internal hosting deployments, Microsoft recommends not exceeding 20,000 transport rules per forest. A functional issue was observed in test in which the transport service became inoperable at levels of scale beyond 20,000 transport rules. Given this limitation, if customer requirements involve one or more transport rules per tenant, then the number of transport rules would limit the number of tenants that could be hosted in a given Exchange organization (or Active Directory forest).

## Accepted Domains

Historically, the number of configured accepted domains within an Exchange organization has been a scalability limit for large-scale hosting deployments. As mentioned in the Tenant Scalability section, each tenant created during scalability testing included an accepted domain. Testing demonstrated that 50,000 accepted domains did not result in any functional issues with the Exchange product. Customers who want to exceed 50,000 tenants should plan to perform pre-deployment validation and stage deployment so that any issues resulting from scalability of accepted domains can be addressed with

minimal impact to the deployment or migration process. It's possible that an extremely high number of accepted domains could limit the number of tenants that could be hosted in a given Exchange organization (or Active Directory forest).

## Public Folders and OAB Distribution

The tenant count mentioned above makes a major assumption that public folder OAB distribution is not being used, which has the side effect of requiring Microsoft Office Outlook 2007 or Outlook 2010 for customers who wish to support Outlook client access in their hosted service offering. Enabling public folder OAB distribution will result in the per-forest tenant count dropping quite drastically to a value that will be determined by a limit in the Exchange system as well as the level of tenant churn occurring within the forest. Given that the true limit isn't very deterministic (because multiple variables can influence it), the Microsoft product group has elected to not support public folder OAB distribution for multi-tenant hosted deployments, as discussed in the [Multi-Tenancy and Hosting Guidance for Exchange Server 2010 SP2](#) white paper. When a public folder-distributed OAB is created, it consumes a replication ID (REPLID). This is an identifier of which there are a fixed number that can be configured in a given Exchange organization. The fixed number is approximately 32,000. REPLIDs are consumed for other purposes as well, but in a multi-tenant hosted Exchange organization, the number of OABs will be the largest consumer of these by far. Once a REPLID has been consumed, it will never be returned to the pool for future use. To receive support from Microsoft, public folder OAB distribution can't be used in a multi-tenant hosted deployment.

## General Recommendations for High-Scale Deployments

Microsoft provides generic performance and scalability guidance for Exchange 2010. For more information, see [Performance and Scalability](#) in the Exchange Server TechCenter Library.

The sizing process for high-scale hosted deployments requires careful attention to domain controller scalability because hosted deployments tend to be associated with extremely large directories. The test results quoted later in this document provide a sense of Active Directory database growth as the number of test tenants increased. During the test process, an observation was made that the directory size for the Exchange 2010 SP2 deployment was *significantly* smaller than the directory size for an equivalent deployment with the Exchange 2010 /hosting option.

Sizing should be performed with an understanding of user concurrency. Rather than sizing for the number of provisioned tenants and users, the deployment should be sized for the maximum concurrent number of users expected on the system. User concurrency can be measured from an existing hosted Exchange deployment, or if transitioning from another messaging platform, equivalent measurements for the supported messaging protocols can be used as a starting point. Without a clear understanding of user concurrency to base sizing calculations on, the safest option is to size for 100 percent concurrency per the Microsoft recommendation for a typical enterprise deployment.

Also, the deployment should be sized in such a way that additional capacity of a given physical resource can be quickly provisioned to handle unanticipated increases in user demand. Microsoft recommends deploying a significant amount of extra capacity during the initial deployment as a buffer to handle

these unanticipated increases in user demand and has found that this is necessary to ensure high availability and acceptable performance of a hosted solution given the dynamic nature of the customer base. As with any high-scale deployment, Microsoft recommends lab testing prior to production deployment to ensure that the specified hardware is capable of meeting the workload requirements. Given the extreme scale of hosted deployments, it may be difficult or impossible to perform lab validation of an entire forest simultaneously, so Microsoft recommends that the design include a “building block” or “scale unit” approach in which a subset of the hardware can be validated as a unit and multiple units can be put together to build the forest.

## Active Directory Optimizations

Large-scale hosted Exchange deployments may see a significant volume of authentication requests at the Client Access server (CAS) role. Two tuning options are available to help ensure that the authentication process works properly in this type of deployment.

When client protocols use NTLM or basic authentication, the Netlogon service on the server that the credentials were presented to will be responsible for validating the credentials via an Active Directory domain controller. A limited number of threads are available for concurrently processing these authentication requests within the Netlogon service. The number of threads is adjustable via the MaxConcurrentApi registry key, as described in the Microsoft Knowledge Base article, 975363 -- [You are intermittently prompted for credentials or experience time-outs when you connect to Authenticated Services](#). The MaxConcurrentApi setting may be increased to handle additional concurrent authentication load, and the impact of tuning this key may be monitored via the Netlogon performance counters described in the Knowledge Base article.

Kerberos authentication may also be tuned via application of the hotfix described in Knowledge Base article 2545833 -- [Slow performance occurs when many user authentication requests are handled in Windows Server 2008 R2](#).

## Exchange Server 2010 SP2 Multi-Tenant Scalability Testing

### Goals

The principal goal was to gradually provision up to 50,000 tenant organizations, each with a unique address book policy assigned to its users, and then determine if there were any performance bottlenecks, scalability limitations, or other impacts to client functionality at that scale. Functionality tests were carried out at defined scale points along the way as the number of tenant organizations grew.

In this testing, “tenant” is synonymous with “customer” (for example, when we refer to 1,000 tenant organizations, it’s the same as referring to 1,000 customer organizations). Having said this, it’s important to note that tenants who are configured with multiple accepted domains or transport rules impact the overall sizing of the Exchange deployment differently than tenants who are configured with a single accepted domain, as discussed earlier in this document.

Another goal was to capture lessons learned and best practices resulting from the testing effort.

## Testing Methodology

The test team deployed a virtualized hosted lab environment of Exchange 2010 SP2 (Enterprise version). Tenant organizations were gradually provisioned via a custom Exchange Management Shell script. Each tenant organization included the following:

- A tenant organizational unit in Active Directory.
- Two global security groups.  
*Note: Global security groups were used to apply permissions to Active Directory objects outside of the Exchange system. As a best practice, universal groups should be used with Exchange 2010 as they can be managed using the Exchange management tools.*
- An accepted SMTP domain.
- A global address list (GAL), filtered by tenant [based on CustomAttribute1].
- An “All Rooms” address list, filtered both by tenant (as above) and a RecipientDisplayType of “ConferenceRoomMailbox”.
- An “All Users” address list, filtered both by tenant (as above) and ObjectClass=user.
- An “All Contacts” address list, filtered both by tenant (as above) and ObjectClass=contact.
- An “All Groups” address list, filtered both by tenant (as above) and ObjectClass=group.
- An offline address book (OAB), based on the tenant organization’s GAL. OABs were configured to be distributed by all CAS servers in the environment and were ‘homed’ in round-robin fashion across the two available mailbox servers. Finally, the OAB update schedule was set to “Never,” which is a best practice for large-scale hosted deployments. Rather than updating OABs on a scheduled basis (the default behavior within the Exchange system), OABs should be updated by the provisioning system on an as-needed basis to reduce the impact of OAB rebuild and distribution operations.
- An address book policy, with the entire GAL, address lists, and OAB linked to it.
- Three mailbox-enabled users, with their CustomAttribute1 set to equal the tenant name. Mailboxes were evenly distributed across two mailbox databases in a database availability group (DAG).
- Two mail-enabled contacts, with their CustomAttribute1 set to equal the tenant name.

After all the above had been provisioned, the tenant OAB was updated (via the Update-OfflineAddressBook cmdlet), and a new transport rule was created, based on a “SubjectContainsWords” filter.

Break points were defined at the following tenant organization counts: 1,000; 5,000; 10,000; 20,000; 25,000; 30,000; 40,000; and 50,000. At each load point, the following information was gathered:

- Average end-to-end time to fully provision a tenant organization
- Average time for each provisioning cmdlet to run

At each break point, the following functionality tests were run:

- **Outlook Web App Test:** Login to a new tenant user's mailbox (for example, USER01@Tenant5001.com) using Outlook Web App
  - Open the Address Book feature in Outlook Web App.
    - Verify that only those address books assigned by the tenant's ABP are visible.
    - Verify each of those address lists are populated as expected.
  - Send an e-mail to another user in the same tenant (for example, user02@tenant5001.com).
- **Outlook Test:** Create an Outlook 2010 user profile for another user in the same tenant (for example, user02@tenant5001.com)
  - Verify that profile creation succeeds.
  - Open the mailbox in Outlook 2010.
  - Open the Address Book feature in Outlook.
    - Verify that only those address books assigned by the tenant's ABP are visible.
    - Verify each of those address lists are populated as expected.
  - Verify that the Outlook client can download the tenant OAB.
  - Switch to 'offline' mode and verify that the OAB is populated correctly.
  - Verify that the e-mail sent from Outlook Web App (in the previous test) has arrived.
  - Reply to the e-mail from the previous test.
  - Go back to the Outlook Web App client, and verify that the reply arrived.

## Scale Testing Lab Setup

This section provides an overview of the number and types of servers in the scale testing lab, and shows how they were distributed to Microsoft Hyper-V root servers.

### Distribution of Virtual Server Images on Hyper-V Root Servers

The following diagram shows the distribution of the virtual server guest images on the Hyper-V root servers.



Distribution of the virtual server guest images on the Hyper-V root servers

### Configuration of Virtual Server Images (RAM, CPU, Disk)

Table 1 shows the hardware configuration of both the Hyper-V root (physical) servers, as well as the RAM and CPU allotments to the Hyper-V virtual server guest images. Don't use this table as a reference architecture; it's simply a definition of a hypothetical test topology. The CPU and memory sizing represented here isn't designed to match the needs of any specific production workload. To correctly size a deployment, use the Exchange 2010 sizing and capacity planning process described in [Performance and Scalability](#).

Hyper-V physical host	Guest virtual server name	Guest cores	Guest RAM
<b>Hyper-V Host01</b>	<b>AD01</b>	<b>4</b>	<b>16 GB</b>
<i>8 Cores, 2.27 GHz Intel Xeon</i>	<b>HUB01</b>	<b>2</b>	<b>8 GB</b>
<i>48 GB RAM</i>			
<i>33-drive RAID 6 Array</i>			
<b>Hyper-V Host02</b>	<b>AD02</b>	<b>4</b>	<b>16 GB</b>
<i>8 Cores, 2.27 GHz Intel Xeon</i>	<b>HUB02</b>	<b>2</b>	<b>8 GB</b>
<i>48 GB RAM</i>			
<i>33-drive RAID 6 Array</i>			
<b>Hyper-V Host03</b>	<b>CAS01</b>	<b>4</b>	<b>12 GB</b>
<i>16 Cores, 2.27 GHz Intel Xeon</i>	<b>MBX01</b>	<b>4</b>	<b>12 GB</b>
<i>32 GB RAM</i>			
<i>10-drive RAID 5 Array</i>			
<b>Hyper-V Host04</b>	<b>CAS02</b>	<b>4</b>	<b>12 GB</b>
<i>16 Cores, 2.27 GHz Intel Xeon</i>	<b>MBX02</b>	<b>4</b>	<b>12 GB</b>
<i>32 GB RAM</i>			
<i>10-drive RAID 5 Array</i>			

Hyper-V physical host	Guest virtual server name	Guest cores	Guest RAM
Hyper-V Host05	AD03	4	16 GB
16 Cores, 2.27 GHz Intel Xeon	Win7Client01	2	2 GB
32 GB RAM	Win7Client02	2	2 GB
10-drive RAID 5 Array			

Table 1 - Hardware configuration of Hyper-V root (physical) servers with guest image allotments

## Detailed Scale Testing Results

### Test Results at Measurement Points

This section presents the performance benchmarks at each phase of scale testing.

The performance benchmarks listed in Table 2 were captured during maximum provisioning load; in other words, tenant organizations (including mailbox-enabled users and mail-enabled contacts) were being created as quickly as possible. All throttling was disabled.

Total tenants provisioned	Total transport rules created	Average time to provision tenant	Outlook Web App test	Outlook test
1,000	1,000	8 seconds	Success	Success
5,000	5,000	12 seconds	Success	Success
10,000	10,000	20 seconds	Success	Success
20,000	20,000	40 seconds	Success	Success
25,000	25,000	50 seconds	Fail*	Fail*
30,000	20,000	60 seconds	Success	Success
40,000	20,000	80 seconds	Success	Success
50,000	20,000	100 seconds	Success	Success

Table 2 - Performance benchmarks at each phase of scale testing

\* After the test environment had scaled up to 25,000 tenants (including 25,000 transport rules), the transport service failed on both hub servers. The service could not be restarted successfully due to a scalability issue with the number of transport rules configured in the system. At this point in the testing, the total number of transport rules was reduced to 20,000 and testing continued with no additional transport rule creation.

### Results at Scale Target

After reaching 50,000 tenant organizations (and 150,000 total mailboxes), the environment was stable and demonstrated acceptable performance. The Outlook Web App and Outlook tests succeeded with no issues. No scalability or performance issues with address book policies were encountered. The only serious issue identified was with total number of transport rules when the value exceeded the recommended limit of 20,000, as described above.

### Validation with Simulated Client Load

After target scale (50,000 tenant organizations) had been reached, load testing was performed against a single virtualized Client Access server to ensure that the Address Book Policies feature remained functional with client load at target scale.

Load testing consisted of Outlook Web App load, which was generated using the Microsoft Exchange Load Generator (LoadGen) tool. The first test pass of Outlook Web App load consisted of 500 concurrent Outlook Web App sessions for 8 hours, running the Outlook Web App 2010 Enterprise script included with the LoadGen tool. The second test pass of Outlook Web App load consisted of 2,000 concurrent Outlook Web App sessions for 8 hours, again running the Outlook Web App 2010 Enterprise script.

Performance was evaluated after the load testing was complete, and no abnormal issues were found in the system.

The load simulation wasn't designed to simulate realistic client load at scale. It was simply designed to ensure that the system remained functional with a moderate amount of client load after the system had been provisioned to the target level of tenant scale.

### Active Directory Sizing Observations

As the number of tenant organizations grows, the number of objects stored in Active Directory grows. In turn, this increase results in an increase in disk space consumed by the Active Directory database file (NTDS.DIT), as shown in Table 3. To assist with capacity planning, the NTDS.DIT size was captured during this scalability testing process. The actual NTDS.DIT size of a given Exchange deployment may vary greatly depending on total number of objects, utilization of various attributes (for example, storage of certificates), deployment of third party tools and/or additional Microsoft products which use the Active Directory database as a configuration store, etc.

Number of tenant organizations	Size of NTDS.DIT in Exchange 2010 SP2 Enterprise
1,000	153MB
5,000	614MB
10,000	1.15GB
20,000	2.3GB
30,000	3.5GB
40,000	4.6GB
50,000	5.9GB

Table 3 - Active Directory sizing observations

As the number of users or other recipient objects in the deployment grows, the Active Directory database will grow as well. The total number of objects stored in the Active Directory database may be a limiting factor on the scalability of a given multi-tenant Exchange forest because the size of the database may impact the time required to on-board new domain controllers, backup domain controller databases, and may also impact the ability of Active Directory to successfully replicate changes throughout the Active Directory topology. For more information about Active Directory scalability limits, see [Active Directory Maximum Limits - Scalability](#). Exchange forests which will be designed to grow to extremely large user counts will require additional pre-deployment testing to ensure that the design is achievable given anticipated rate of Active Directory object churn, requirements for time to on-board new domain controllers, requirements for time to recover from a domain controller hardware failure, and other specific deployment requirements.

## Summary

Building a large-scale multi-tenant Exchange 2010 deployment is an extremely complex undertaking. Extreme care should be taken during the design, pre-production test, and deployment phases to ensure that all aspects of the system scalability have been considered and are being closely monitored. No large scale deployments of this type are the same—every deployment has customizations that impact how the system will function as the number of tenants and users increases. Therefore, testing, monitoring, capacity planning, and risk mitigation are all critical aspects of a successful multi-tenant Exchange deployment.

This document provides some high-level scalability guidance and deployment recommendations. We welcome your feedback and comments on this document and will incorporate them into future versions if required. If you have feedback, please send it to [ExHostingFeedback@microsoft.com](mailto:ExHostingFeedback@microsoft.com).

## Legal Notice

This document is provided “as-is”. Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it.

Some examples depicted herein are provided for illustration only and are fictitious. No real association or connection is intended or should be inferred.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

© 2011 Microsoft Corporation. All rights reserved.

Microsoft, MS-DOS, Windows, Windows Mobile, Windows Server, Active Directory, ActiveSync, Forefront, Outlook, and SharePoint are trademarks of the Microsoft group of companies.

All other trademarks are property of their respective owners.