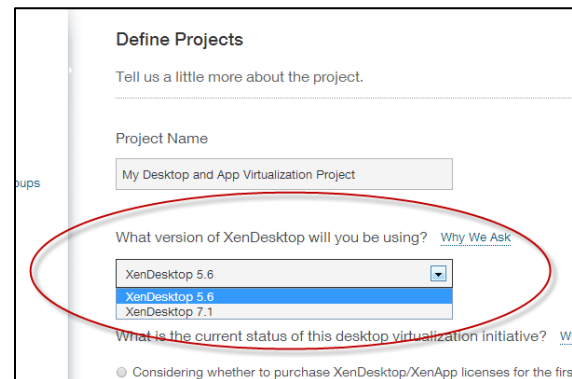


Frequently Asked Questions

1) Why is my project not showing XenDesktop 7 results?

Project Accelerator has been upgraded to support XenDesktop 7. Because older projects and existing implementations might be based on XenApp 6.x (IMA) or XenDesktop 5.6 we need you to tell us what version of the product you will be using. The scalability, design constraints, and Citrix Consulting recommendations are different for those versions of the XenDesktop and XenApp platforms. You can set your existing project to use XenDesktop 7 by opening the “Define Project page” (Step 2 in Assess) and picking the value ‘XenDesktop 7.1’ (works for 7.5 and 7.6 too) for the “What version of XenDesktop will you be using” question.



The screenshot shows a web form titled "Define Projects". It includes a text input field for "Project Name" with the value "My Desktop and App Virtualization Project". Below it is a dropdown menu for "What version of XenDesktop will you be using?". The dropdown is open, showing three options: "XenDesktop 5.6", "XenDesktop 7.1", and "XenDesktop 7.1". The "XenDesktop 7.1" option is highlighted in blue. A red oval is drawn around the dropdown menu. Below the dropdown is a radio button for "Considering whether to purchase XenDesktop/XenApp licenses for the first time".

2) How do I login?

Project Accelerator requires a “My Citrix” credential. Its free, you can [create one here](#). If you are a Citrite your “My Citrix” credential is the same as the Citrite credentials you use for email and other employee applications.

3) How do I “Customize my Design”?

“Customize my Design” is now available as part the Design phase of Project Accelerator! You can change everything from FlexCast model to application delivery to resource allocations, resulting in a design tailored to your organization’s environment and end user preferences. There are thirty-four decisions organized in the five layers Citrix recommends for representing a XenDesktop architecture.

4) Why should I trust these numbers?

You should review the recommendations with your Citrix representative or trusted partner. Nevertheless these numbers are based on detailed logic developed by Citrix Consulting’s experts in desktop virtualization Assessment and Design. The logic and derived results are reviewed and improved based on input from customer engagements, lab tests, and technical studies. The numbers presented are based on what Citrix Consulting would recommend if you gave them the same information.

5) What do the Original and Customized buttons do?

On both the Sizing and Deployment Plan and Architecture Diagram, there are buttons marked “Original” and “Customized”. When you select “Original” you see the Citrix Recommended values based on the Organization, User Group and Application questions you answered during the Assess phase. If you change answers in Assess, the “Original” values will be updated. The “Customized” button shows the results of your Design decisions. At any time, you can see how the hardware sizing and architecture depart from the Citrix recommendations using the Customized v. Original “toggle” on each output. The Export to PDF button will create a PDF version of the active version of the output.

6) What about local storage?

We are considering the addition of logic to allow for local storage in a future release.

7) Why are my infrastructure sizing estimates higher than I expected?

Infrastructure calculations include XenDesktop & XenApp controllers, Provisioning Services servers, StoreFront servers, XenClient Synchronizers, database and licensing servers, and XenApp servers for application hosting as appropriate for the recommended environment. For the initial recommendation, we are designing the database as a 3-node mirror, N+1 level of fault tolerance for XenDesktop, XenApp, Provisioning Services, and StoreFront servers.

Even though XenApp servers for application hosting are contained within the Infrastructure line within the Sizing Plan, these servers are depicted in the Architecture Diagram as App Hosts and App VMs within the XenApp Hosts box in the Hardware Layer. Storage space and IOPS are also impacted by the XenApp servers for application hosting.

8) When should I use AppDNA or Lakeside software to help with application virtualization?

Citrix AppDNA software automates application compatibility testing, remediation and preparation, while Lakeside FastTrack shows users’ usage behavior and resource demands. The systems can be used separately or together: AppDNA software consumes data from Lakeside and combines it with AppDNA application information to create a clear picture of which user groups can be easily deployed first and which

will require more migration work. As you step through Project Accelerator, free Citrix AppDNA and Lakeside trials are made available.

9) How are FlexCast model recommendations calculated?

FlexCast model recommendations are provided using Citrix Consulting leading practices, a collection of our knowledge of the technical abilities and benefits of each FlexCast model. We also want to dissuade you from making poor choices and so on occasion we remove possible choices.

As you enter business priorities, skills and user group requirements into the Assessment wizard, we use our knowledge of what will work best to determine the optimal and allowed FlexCast model for each user group.

Some models may be invalidated completely meaning they should not be used for this user group. Those FlexCast models will be removed from the choice list in Design. This does not mean it is technically impossible but it does mean that the end-user or business requirements will not be met with those FlexCast models.

For example, your organization has a BYOD initiative and also want to increase security (business priorities); they also do not have any XenDesktop or XenApp knowledge (skillset). The Marketing group needs “Complete” personalization, “Medium” security, “Remote” mobility, some resource intensive apps and will have a “Heavy” workload. We would suggest Assigned VDI, but Local VM and other models would be appropriate. However, Hosted Shared and On-Demand Apps will not support Complete Personalization and those FlexCast models will be removed from the choice list in Design.

10) How are Hardware estimates for Compute calculated?

Compute estimates for the Hosts in the Datacenter are shown as Cores and RAM. The constants we use to determine the recommended RAM and Cores are based on Citrix Consulting’s experience in the field and lab testing that simulates production environments.

We use different calculations for the user groups and infrastructure items in the Sizing and Deployment Plan. Each of the two calculations for a user group are based FlexCast model, defined Workload and the Windows Operating System version. User groups using FlexCast models of Local VM, Streamed VHD and Remote PC do not require any compute in the datacenter unless they use applications shared by XenApp.

Cores:

To calculate the number of cores for a user group, we use the following calculation:

$$\begin{aligned} & Roundup \left(\frac{\text{NumberOfUsers}}{\text{UsersPerCore}} \right) + Roundup \left(\frac{\text{NumberOfXenAppDedicatedUsers}}{\text{UsersPerCore (HS, normal, 2012R2)}} \right) \\ & + Roundup \left(\frac{\text{NumberOfXenAppSharedUsers}}{\text{UsersPerCore (HS, light, 2012R2)}} \right) \end{aligned}$$

The constant values for users per core in Project Accelerator are:

	Operating System	Light Workload	Normal Workload	Heavy Workload
Pooled VDI	Windows 7	13	10	5
	8	15	11	6
	8.1	15	11	6
Assigned VDI	7	11	8	4
	8	13	9	5
	8.1	13	9	5
Hosted Shared and On-Demand Apps (ODA)	Win Server 2008 R2	21	14	7
	2012 R2	21	14	7

Table 1 - Desktop Users per Core

Note that the scalability here and elsewhere for both types of XenApp delivery models, Hosted Shared and On-Demand Apps (ODA) are the same. For example, if we have 500 users, with a normal workload and no usage of shared or dedicated apps on a hosted shared desktop running Windows Server 2012 R2, we get the following:

$$Roundup \left(\frac{500 \text{ users}}{14 \text{ user per core}} \right) = 35.714 \text{ Cores} \approx 36 \text{ Cores}$$

RAM:

To calculate the amount of RAM for a user group using Pooled or Assigned VDI, we use the following calculation:

$$\begin{aligned} & (\text{NumberOfUsers}) * (\text{RAMPerDesktop}) \\ & + (\text{Shared\&Dedicated Application Virtual Machines}) * (\text{RAMperVM}) \end{aligned}$$

To calculate the amount of RAM for a Hosted Shared user group we use the following calculation:

$$(NumberOfVirtualMachines) * (RAMPerVM) \\ + (Shared\&Dedicated\ Application\ Virtual\ Machines) * (RAMperVM)$$

The RAM constants we currently use in Project Accelerator are:

	Operating System	Light Workload	Normal Workload	Heavy Workload
Pooled VDI	7	1.5	2	4
	8	1.5	2	4
	8.1	1.5	2	4
Assigned VDI	7	2	3	4
	8	2	3	4
	8.1	2	3	4
Hosted Shared	2008 R2	12	12	12
	2012	24	24	24
	2012 R2	24	24	24

Table 2 - RAM per workload

For example, if a user group requires 100 Pooled virtual machines on Windows 8.1 with a heavy workload requirement but no shared or dedicated applications, we get the following:

$$RAM = (100 \text{ virtual machines}) * (4 \text{ GB per VM}) = 400 \text{ GB}$$

11) How are storage sizing estimates calculated?

Storage estimates for the datacenter are broken into two parts: Space in gigabytes (GB) and Input-Output operations per second (IOPS). User groups using FlexCast models of Local VM, Streamed VHD and Remote PC do not require any datacenter storage space or IOPS unless they use applications shared by XenApp.

Space:

To calculate the amount of storage for a user group, we use the following calculation:

$$\text{Roundup} (\text{TotalCache} + \text{PersonalvDisksize} + \text{WorkerGroupCacheSize})$$

The default values for the amount of storage for a medium workload are:

	Hosted Shared/ODA	Pooled VDI	Assigned VDI
Personal vDisk Size	n/a	n/a	10GB
Desktop Image Size	60 GB	35 GB	35 GB
Cache Size (medium workload)	25/40 (2008/2012)	7 GB	7 GB
MCS Thin Provisioning	25%	25%	75%

Table 3 - Storage Constants

For example, if we have 100 Pooled VDI virtual machines for 100 users with a Medium workload which are delivered by Machine Creation Services (for which we assume Thin Provisioning and an identity disk sizing 16MB) and no shared or dedicated applications being delivered by XenApp, we get the following space calculation:

$$\text{Space} = \text{RoundUp} \left(100 * \left((35 * 0.75) + 0.015625 \right) \right) + (100 * 10) + (0) = 3627 \text{ GB}$$

There is a different calculation used for calculating infrastructure storage and master image size.

Average IOPS:

To calculate the amount of steady-state average IOPS for a user group, we consider the FlexCast model, need for shared or dedicated applications, imaging solution, and operating system. Here is the calculation we use:

$$\begin{aligned} & (\text{NumberOfUsers}) * (\text{SteadyStateIOPSPerUser}) + (\text{Number of Shared Users} * \\ & \text{IOPS Constant for HS Light, 2012R2}) + \\ & (\text{Number of Dedicated Users} * \text{IOPS Constant for HS Normal, 2012R2}) \end{aligned}$$

The default values for the number of average IOPS per workload are:

FlexCast	Operating System	Light			Medium			Heavy		
	Imaging Solution→ Operating System ↓	Manual	MCS	PVS	Manual	MCS	PVS	Manual	MCS	PVS
	7	7	7	5	13	13	10	26	26	20
	8	7	7	5	13	13	10	26	26	20
	8.1	7	7	5	13	13	10	26	26	20
	7	NA	5	4	NA	12	10	NA	26	20
	8	NA	5	4	NA	12	10	NA	26	20
	8.1	NA	5	4	NA	12	10	NA	26	20
Hosted Shared and On-Demand Apps	2008 R2	3	3	2	6	6	4	12	12	8
	2012	5	5	3	9	9	6	17	17	12
	2012 R2	5	5	3	9	9	6	17	17	12

Table 4 - IOPS per workload

For example, if a user group has 500 light users using Pooled VDI and Machine Creation Services and no shared or dedicated applications are needed and the Steady state IOPS decision in Design is left at the default value, we get the following

$$(500 * 7) = 3,500 \text{ Average IOPS}$$

12) Why are the IOPS numbers missing from the Infrastructure row on the Sizing Plan?

The Infrastructure components (identified as the Control Layer within the Architecture Diagram) do require IOPS; however, these systems do not incur significant IOPS relative to the other parts of the design so we have not implemented a calculation for infrastructure IOPS.

13) How is the recommended deployment order determined?

We determine the deployment order by ranking each user group against how well they align to the business priorities and how long it will take to implement.

Based on the rankings entered into the Assess phase, we are able to link business priorities with each user group's requirements. For example, if Cost is the highest business priority, user groups with workload requirements of "Light" will have a higher business impact than those users with workload requirements of "Heavy" or "Normal". We look at all business requirements and rank each user group based on who aligns closest with the stated priorities.

Next, we look at how long each user group's solution will take to implement. This calculation centers around the existing skillsets within the organization. We also look at all technical skillsets and rank each user group based on which group should take the shortest amount of time to implement.

In order to come up with a final ranking we sum the rankings from Business Impact and Time to Implement. This provides us with the recommended deployment order.

14) Is redundancy, disaster recovery or high-availability factored into Project Accelerator recommendations?

The initial recommendation is based on a failover model of N+1 for infrastructure components, meaning we add an extra server into our calculations. For example, if our calculations say we need 4 virtual machines to host Provisioning Services, we will recommend 5 to provide enough capacity in the event of a single virtual machine becoming unavailable.

We are not calculating any level of failover or Disaster Recovery in Project Accelerator.

15) Why do Remote PC, Local VM and Streamed VHD have no resources allocated in the Sizing Plan?

These three FlexCast models rely on client-side resources (cores, memory, storage space and IOPS). The sizing plan is focused on identifying how many server-side resources are required to support the user group. So unless the user groups using these FlexCast models need shared or dedicated applications their user group lines will show n/a.

Even though these FlexCast models do not impact resources within the user group, they do have an impact at the control layer. These three FlexCast models require appropriate controller(s), which are reflected in the Control Layer of the Architecture diagram and the Infrastructure line in the Sizing Plan.

Accelerator Plus Known Issues (Plus Partners only)

Project Accelerator Plus is a special offering of Project Accelerator we are developing for Citrix Solution Advisors. It is only current only available to a few partners who have elected to participate in the Beta program. This portion of the FAQ refers to those features and known issues with their use.

16) Two Dynamic Deliverables

Currently only two dynamic deliverables are available. So the menu for the other templates only allows for a download. We hope to add more dynamic deliverables in the future.

17) Images are blank in emailed Dynamic Deliverables documents

If a dynamic deliverable is downloaded to a Windows computer and saved as a .docx file format and then emailed to a recipient using an MAC version of Office or vice versa, the images in the document may be stripped. This is due to different interpretations of HTML by Office on each platform. If you save the file as a .DOC (2007 version) file, the images can be viewed across platforms.

18) Page numbering

Page numbering in the dynamic deliverables does not always populate correctly in the footer of the document or in the Table of Contents. If it does not, it can be fixed by correcting the page numbers to continue from the previous Word document section. To do this correction, navigate to the footer of the document and then use the Format Page numbers item within the Page number section of the Header/Footer Tools special menu to correct the page numbers to continue from the previous section. Then use the context menu (right-click) on the Table of Contents and select the appropriate menu to update the page numbers there.

19) Template auto-replace

The following fields are not available yet for automatic replacement via the advanced properties function: consultant1, consultant2, address, number, email1, email2 do not appear in advanced properties but can be replaced by adding them manually and following the same instructions as updating the other fields. Also the Consulting Team field on the Title Page does not auto-update correctly. Please correct it manually.

20) No Dynamic Deliverables for XenDesktop 5.6 projects

If you have an old project or have selected XenDesktop 5.6 on the project page in Assess, you will not be able to populate the Dynamic Deliverables. Go back to your project page for the Project Accelerator project you've created and select XenDesktop 7, to be able to use the menu on the Conceptual Design and Use-Case Analysis Dynamic Deliverable.

21) Minor known issues

- On Safari, IE, and Firefox, users cannot generate the Conceptual Design dynamic deliverable for a project with 7 or more datacenters.
- When not connected to the Internet, the architecture diagram may not be visible.
- If Word prompts for a file type, select "Save-as Web"
- The templates and dynamic deliverables that indicate they support XenDesktop 7.1 also support XenDesktop 7.5 and XenDesktop 7.6
- Some templates have an old, static version of the architecture diagram. Replace with the one from Project Accelerator or feel free to generate your own using other tools.