

Inside Citrix chapter twenty seven – The one with Appendix A - FMA facts

1. There is also a Linux-based VDA
2. Active Directory is required for the authentication and authorisation of users in a Citrix environment. This includes DNS.
3. Your Delivery Controllers can be considered as the heart of your FMA deployment.
4. Your environment is as strong as its weakest link. Make sure to apply the 'one is none' rule wherever and whenever it makes sense.
5. Prior to XenDesktop 7 the VDA was referred to as the Virtual Desktop Agent, while today we know it as the Virtual Delivery Agent, a subtle difference.
6. You can configure multiple Machine Catalogs with different desktop and server Operating Systems within the same environment / Site.
7. If a VDA is unable to register itself with a Delivery Controller or communication between the VDA and the Delivery Controller fails for any reason, the machine will stay in an unregistered state and won't be directly accessible or manageable through one of your Delivery Controllers.
8. There is a separate HDX 3D Pro VDA for use with GPU acceleration for example. This type of VDA enables you to make use of hardware acceleration, including 3D professional graphics applications based on OpenGL and DirectX. (The standard VDA supports GPU acceleration of DirectX only.). It can be selected during VDA installation. Resources can either be assigned on a one to one basis (Passthrough) or shared amongst multiple VMs (vGPU).
9. While XenDesktop and XenApp both support Web Interface (EOL June 2018) Citrix recommends using StoreFront for new as well as existing deployments. It is built for the future and as such has a whole bunch of additional features not available in Web Interface.
10. Note how I mention user authentication and user validation. There is a distinct difference. Authentication is to make sure that somebody is who he or she claims to be. Verification is done to find out which resources are assigned (permissions) to the user.
11. Note that besides the Receiver for Web approach, where users log into StoreFront by means of a web page, you can also configure your Citrix Receiver in self-service mode. This way your users will be able to subscribe to their resources directly from the local Citrix Receiver interface. See the 'The Citrix Receiver' section for some more detailed information.
12. Besides using Keywords, as of Citrix Receiver 4.2.100 you can also integrate application and desktop short cuts into your user's Start menus or put them onto their desktops, with no resource subscription needed.
13. Going forward, StoreFront multi-site configurations will be a lot easier to configure and implement. Most functionality will be built into the Graphical User Interface of StoreFront.
14. A XenApp Farm (6.5) or XenDesktop / XenApp Site (7.x) is also referred to as a 'Deployment' by Citrix. Especially if you spend some time on their E-docs pages you, will see this term a lot.
15. We can use the Optimal NetScaler Gateway routing feature to route the user's ICA traffic through the NetScaler most applicable (the one connecting them to their

- XenDesktop Site in the case of a multi-site deployment) to the user, even if the initial connection was made through another NetScaler.
16. By default, StoreFront will use your internal services URL as an internal resolvable Beacon point and it will use Citrix.com as the external Beacon point. But you can change them to whatever you like. Just make sure that your internal Beacon is not resolvable externally.
 17. Non-Platinum-licensed customers can keep and store data for up to 7 days, while a Platinum license allows you to store all data for up to a year, with the default being 90 days.
 18. Connection Leasing is meant to supplement SQL High Availability set-ups.
 19. The Receiver X1 combined with StoreFront will greatly simplify overall management and improve the user experience on multiple levels.
 20. HDX is not a replacement for the ICA protocol. It offers a set of capabilities or technologies that offer a high-definition user experience, which are built on top of the ICA remoting protocol.
 21. While some think that ThinWire is still a relatively new technique, it is not. ThinWire has always been there. It is a core component of the ICA virtual display channel stack (for over twenty years now). That's why they rebranded their latest addition as ThinWire Plus, although it has had several names along the way.
 22. If you want to make use of e-mail-based discovery you will need to use StoreFront.
 23. All, or at least most, of these resource short cut management options were already available with Citrix Receiver Enterprise up to version 3.4, when they killed it. It took up to Citrix Receiver version 4.2 to get this functionality back.
 24. By disabling the SelfServiceMode (it is enabled by default) subscribed-to applications can only be accessed through the Start menu and desktop short cuts. This is also referred to as short cut-only mode.
 25. By default, Studio communicates with the Controller on TCP port 80.
 26. While Studio takes care of most configuration and maintenance tasks, depending on your set-up, it doesn't cover everything. If you are using Provisioning Services, you will still have a second, separate management console. The same applies to Citrix NetScaler.
 27. Do not compare FMA-based Zones (7.x) with IMA-based Zones (6.5). There are some distinct differences between the two. Make sure to check out the table on page 101.
 28. If the RRT to and from a satellite Zone is near or above 250 ms, a separate Site deployment, including an SQL HA set-up, is advised.
 29. If you want to limit the number of brokering requests originating from a satellite Zone there is a Registry Key, which can be configured for this.
 30. Make sure to check out CTX139382 for a whole bunch of best practices around Director.
 31. As it stands today, the EOL for EdgeSight has been set to 30-June-18, or 24-Aug-2016, depending on if you have a valid software maintenance and/or Subscription Advantage. In that case, the EOM is set to 31-Dec-17 or 24-Feb-2016.
 32. As of version 7.7 Director can be configured to make use of integrated Windows authentication so that domain-joined users gain direct access to Director without re-entering their credentials on the Director logon page.
 33. The actual SCOM web interface can be launched from within Director as well. You will find it on the 'Alerts' page.

34. Director can also be used to monitor and troubleshoot IMA-based architectures in the form of XenApp 6.5. Features include, but are not limited to: Shadow sessions, Machine details pane, HDX panel, Delegated Administration support, and Activity Manager for 6.5.
35. By default, you can only use one type of license within your XenDesktop Site. You either purchase / upload user/device or concurrent: they cannot be mixed. If you require both, you must set up and configure separate Sites, license servers included.
36. The license server uses tables to track user\device license (assignment) information (as described above).
37. The process of assigning licenses to users and/or devices, whether concurrent or not, is also referred to as the checking in and checking out of license tokens.
38. Both XenDesktop and XenApp product licenses must be purchased with Subscription Advantage or Software Maintenance for a minimum of one year from delivery.
39. As soon as a Citrix product enters a grace period, one or several event messages (Windows Event Viewer) might appear. Here you can also see the remaining time left within the grace period.
40. All Session Hosts as part of the IMA are responsible for the checkout and handling of licenses, and thus need to be able to communicate with the license server. Within the FMA this is handled by your Delivery Controller (s).
41. When licenses are allocated they are 'bound' to your license server, which is identified by its local hostname and is Case Sensitive.
42. You can also visit the Citrix Trial Center where you can get limited trial licenses to try out certain products. However, note that some licenses will only be available for registered Citrix partners.
43. Citrix also offers Appliance Maintenance, which provides technical support to diagnose and resolve issues encountered with appliance hardware with the latest upgrades for the software elements of hardware products. Malfunctioning appliances are also replaced under this agreement to minimise customer downtime. Note that all licenses within a programme must be either on call-in support or not – they cannot be mixed! If one desires different support levels, different licensing subscriptions must be used to separate these, as well as separate license servers!
44. Technically speaking, Software Assurance is an upgrade of existing licenses (usually OEM). That's why you cannot have SA on thin clients (there is no existing license to upgrade) and you have to buy VDA license instead.
45. If you are not accessing a Windows desktop OS VM on a server, but from a physical PC, you do not need a Windows VDA license. This also means that VDA licenses do not apply to Citrix XenApp.
46. Software Assurance benefits (either per use or device) allow you to have up to four virtual machines (VDI), or one physical machine running a Windows desktop Operating System.
47. Microsoft RDS licenses are needed in combination with Citrix XenApp, not XenDesktop. And Microsoft VDA licenses are needed in combination with Citrix XenDesktop, not XenApp.
48. Another thing to keep in mind when trying to achieve 'true' cloud based VDI, is that customers will have to provide their own (Windows desktop OS) licenses. A Service Provider is not allowed to sell these.

49. You might have heard about the Nutanix Acropolis Hypervisor. It will soon be available as a Host Connection within XenDesktop as well.
50. Just recently, Citrix introduced the CPX model, which is Citrix's containerised version of NetScaler; mainly used for testing and development use cases. It is still in tech preview at the time of writing.
51. While there is a separate NetScaler Gateway license available, also know that each 'normal' ADC NetScaler (Standard, Enterprise or Platinum license) includes the Gateway functionality by default: no additional licenses needed.
52. The virtual NetScaler (VPX) can handle up to 1500 concurrent ICA connections (supported by Citrix, theoretically it can handle more). If you need more, then you'll have to upgrade and purchase a physical MPX appliance, which, depending on the model, can handle anything ranging from 10,000 to 35,000 concurrent ICA connections at a time.
53. There's a lot of overlap between the two (ADC and Gateway): it basically all comes down to the license you purchase and upload, with the NetScaler Gateway license being the most 'basic' one.
54. A NetScaler SNIP address is probably best compared to a layer 3 routing table entry. Not only does it tell the NetScaler that it has a connection to a specific network, so it is 'known', it also tells it how and where to reach it so that it is able to route network traffic its way.
55. You can also configure a SNIP address as a management IP, instead of, or better said, alongside the NSIP address used to manage your NetScaler.
56. You can configure as many Unified Gateway virtual servers as you like or need.
57. vDisk updates can be automated and scheduled. This feature supports updates detected and delivered from WSUS and SCCM Electronic Software Delivery servers.
58. Be aware that while promoting the version, PVS will actually open up the vDisk and write to it. This it can lead to inconsistencies if you are storing vDisks locally and replication can be complicated. Provisioning Services has its own built-in TFTP server. However, you are free to use whatever you prefer.
59. As an added advantage, using the BDM method will also decrease boot times by around 5 to 10 seconds since we don't have to wait for PXE and TFTP.
60. When vDisks are stored locally on the Provisioning Servers, you will need to implement some sort of replication mechanism so that all PVS servers will be able to offer the exact same vDisks. This can also be done manually from the PVS management console. Recommended automation methods include both DFS-R and Robocopy.
61. The streamed wizard supports the following Hypervisors: XenServer, Hyper-V through SCVMM and ESX through vCenter.
62. Personal vDisks can only be assigned to an desktop Operating System; server OSs are not supported at this time.
63. While I use the term 'provisioning' do not confuse the provisioning of machines with MCS with that of PVS (see previous chapter). In general, provisioning means providing or making something available. A term widely used in a variety of concepts within IT.
64. Today technologies like application layering and containerisation can help us overcome most of these application-related issues; however, the general adoption of these kinds of technologies and products will still take some time.
65. While all services closely interact with and depend on each other, at the same time they are also completely separated from each other. Each service is configured to communicate to the Central Site database using its own individual DB connection string.

- If one service fails, unless they depend directly on each other, it will not affect any or most of the other services.
66. Keep in mind that if you change something for one specific service, like the DB connection string for example, you will have to do this for all of the other FMA services as well.
 67. All FMA services run under the NT AUTHORITY\Network service account. Also, when authenticating to the Central Site database (this is where the Configuration Service plays an important role as well) all services use the local computer account of the machine that they are currently running on.
 68. While it is considered a best practice to keep all Delivery Controllers equally configured, Site services are the exception to the rule, so to speak.
 69. Each FMA service can query the configuration service to look up other services using the listing mentioned earlier. In short, service registration and communication are both reliant on the configuration service. It will also store configuration metadata for all services, relieving Active Directory.
 70. If you would like to refresh the cache of one of the FMA services (remember the five minutes), all you have to do is restart the accompanying Windows Service. The cache (services listing) is retrieved during service start-up.
 71. If you do not configure a Host Connection within Studio, when creating a new Device Catalog, the option to use MCS as a provisioning mechanism will not be available (greyed out). Restarting the Citrix Desktop service on the VDA triggers the registration process and can be used to force re-registration when needed.
 72. As opposed to the Desktop VDA, which has been around for a couple of years now, there is no PortICA service within a Server VDA, it simply does not exist.
 73. Each Terminal Server protocol (like Citrix's ICA) will have a protocol stack instance loaded (a listener stack awaiting a connection request). When installed, the Server VDA basically extends Microsoft's RDS protocol with the ICA/HDX feature set / protocol.
 74. Each service group has a unique identifier, which can be queried using the PowerShell SDK if and when needed.
 75. The ICA protocol originated with Citrix Multiuser, around 1990 / 1991, meaning that the ICA protocol is actually over 25 years of age already.
 76. By default, the ICA protocol uses TCP port 1494. If Session Reliability is enabled a.k.a. the Common Gateway Protocol, or CGP then ICA traffic will be encapsulated through TCP port 2598. Note that any network traces that you might run will also show 2598 instead of 1494.
 77. As a (security) best practice Citrix recommends disabling any virtual channels that are not in use.
 78. As mentioned, there are 32 virtual channels in total; however, Citrix reserves 17 of those. Third-party companies and customers who want to design and implement their own virtual channels are free to use the other ones. These are also referred to as dynamic virtual channels or DVCs.
 79. Other ways to accelerate ICA traffic would include Citrix policies, which can then be applied either per user or per server, or to the whole Site. Implementing a physical accelerator like the Citrix CloudBridge, formerly known as Branch Repeater, is always optional as well.
 80. When not using a CloudBridge appliance, formerly known as Branch Repeater, Session Reliability must be enabled for Multi-Stream ICA to function.

81. When Session Reliability is enabled users will be automatically reconnected as soon as the network connection is reinstated, and they will do so without needing to reauthenticate. Configuring the 'Auto client reconnect authentication' policy to prompt users to reauthenticate can change this behaviour.
82. Remember, Citrix HDX isn't a replacement for the ICA protocol. HDX technologies are meant as an extension and as such operate on top of the ICA protocol.
83. Make sure you check out the HDX policy templates in Studio. There are 6 in total.
84. If you go to YouTube and search for Citrix Framehawk you will find multiple comparison clips of Framehawk vs. other technologies. Guess who comes out on top?
85. Note how I say 'true' application virtualisation. This is because solutions like XenApp are also often referred to as application virtualisation solutions, so it is really a matter of perspective.
86. Published App-V applications can be configured to be launched from the Start menu, through Citrix Receiver, using the locally installed (image) App-V client or from the StoreFront web interface.
87. AppDisks will be available with all XenDesktop / XenDesktop editions, Advanced, Enterprise and Platinum. Note that AppDNA will be for Platinum-licensed customers only.
88. Citrix AppDisks is available as of XenDesktop / XenApp version 7.8
89. Knowing the architecture, the components, the way traffic flows throughout and expected behaviour is the only way to successfully troubleshoot your FMA-based infrastructure.
90. *If you don't enable authentication on the NetScaler's login page the NetScaler will contact StoreFront and the user will be presented (through the NetScaler) with the StoreFront login page (Receiver for Web sites). The user fills in his or her credentials and authentication will be handled by StoreFront.
91. The STA is only used when traffic traverses a NetScaler, so you don't have to worry about the STA service and its tickets when authentication takes place internally. The STA ticket gets generated and sent back after a user launches an application/desktop, and not during the resource enumeration process. It also includes information on the resource to be launched, including the server to launch the application on (load balance).
92. Make sure that the Broker (XML/STA) service on the NetScaler and the StoreFront server is configured identically. The same applies to the load balance/fail over order in which you configure them.
93. When a Delivery Group gets created, two access rules are created and added by default, one for direct connections and one for connections through NetScaler. Using PowerShell we can look at and change these access rules, as we see fit.
94. With XPS, the earlier mentioned print output is already in an XML format and will be sent over to the print spooler service right away. See image on the next page for an overview.
95. Perhaps you are better off using None and Shared mode in production and use Isolated for troubleshooting purposes only, which of course could apply to production as well, only temporarily.
96. As a side note, most thin client devices are based on Linux, as a result they will not be able to locally handle and process the earlier mentioned print jobs. As a result of this, the client printing pathway will only work with Windows-based (fat) client devices.

97. So you see that it's not just one thing, it is everything combined that makes or breaks your print architecture: the type of end points you use, policies configured, including the physical placement of your machines, including printers.
98. If for whatever reason the Citrix (XenApp) server and the print server are unable to communicate with each other, again the client printing pathway will be used (forced) instead.
99. Proper testing will be necessary to ensure that (enough) compression takes place.
100. Once enabled you might want to have a look at the 'Universal print driver usage and preference' policies. You have a bunch of options to select from.
101. Is printing slow? Remember that it isn't just about the bandwidth exclusively. Make sure to check for congestion and latency.
102. When thought through beforehand this phase will probably take you somewhere between 30 to 60 minutes, depending on the number of users and user groups, including the number and the type of applications involved.
103. A high number of IOPS is useless unless latency is low! Even with SSDs which are capable of providing a huge number of IOPS compared to traditional HDDs, latency matters. Latency tells us how long it takes to process a single read or write I/O request.
104. Latency is king: the less you have, the faster your infrastructure will be! Also, there is no standard when it comes to measuring IOPS! There are too many factors influencing overall performance and thus the number of IOPS.
105. Although the average amount of IOPS, or the Steady State, does tell us something, it isn't sufficient. We also need to focus on the peak activity measured between the boot and the Steady State phases and size accordingly.
106. Storage throughput isn't the same as IOPS. When we need to be able to process large amounts of data, bandwidth becomes important: the number of GB/sec that can be processed. Although they do have an overlap, there is a clear difference between the two.
107. If IOPS are limited, try (pre-)booting your machines at night. Also, make sure your users can't reboot the machines themselves.
108. Launching applications will generate high read I/O peaks and initial low writes. Chances are that after users log on they will start, either automatically or manually, their main applications. Again, this is something to take into account, as this will probably cause an application launch storm, although it's usually not recognised as such.
109. By leveraging RAM for writes, a.k.a. RAM Cache with Overflow to Disk in terms of Citrix PVS write cache, we can significantly reduce the number of IOPS needed. In fact, Citrix claims to only need 1 to 2 IOPS per user on a XenApp environment without any complex configurations or hardware replacement.
110. CWC supports both MCS as well as PVS for machine provisioning.
111. Because of Microsoft's licensing restriction with regard to desktop Operating Systems, it is very hard to come up with a true DaaS solution based on an actual desktop OS while keeping costs acceptable. With CWC you can host all of your infrastructural components up in the cloud and leverage your own on-premises VDAs, which can be VDI and 100% desktop OS-based deployments. While not exactly the same, it comes close to a desktop OS-based DaaS (private cloud) offering.
112. The Cloud Connector is what your VDAs will point to and use as a broker, instead of a Delivery Controller when compared to an on-premises deployment.
113. You will need to set up at least two Cloud Connectors per Resource Location to achieve HA. You won't have to configure load balancing in any way for these two Cloud

Connectors. CWC will send requests and data to one of the two Connectors, and if it gets too busy or stops responding, the data will be sent over to the idle Connector, or the load will be spread amongst the two.

114. The customer's metadata will always be stored separately for each tenant, and secured with unique credentials.
115. Citrix offers out-of-the-box blueprints for XenDesktop, XenApp, XenMobile, NetScaler and the Workspace Suite.
116. As I mentioned earlier, E2E relies heavily on sponsors: they are the ones who make it all possible. The E2EVC team doesn't make any money from the event: everything that comes in is invested right back into the event itself.