EuroCloud

Whitepaper

Cloud Computing Guide

State and future of Application Performance Measurement in today's complex Cloud environments

11

Table of Contents

1	Management Summary	4
2	Application Performance Management (APM) Today	5
3	The future of APM	7
4	Today's and future VDI Cloud monitoring solutions	9
5	Benchmarking, Load- / Stress-Testing and Regression Test	17
6	More detailed view based on Desktop as a Service (DaaS)	17
7	Legal Notice	26
8	Author	27

Credits

EuroCloud.Austria - Verein zur Förderung von Cloud Computing
Museumstrasse 5/14
1070 Vienna, Austria
E-mail: info@eurocloud.at
Web: http://www.eurocloud.at
Seat of the Association: Vienna
Copyright: EuroCloud.Austria 2014



Dear Reader; Cloud, mobile and social megatrends have resulted in unprecedented levels of complexity in today's IT environments. As a result, more components of the application delivery chain are obscured from IT and line of business owners. Performance visibility and greater operational intelligence should be paramount to all organisations amid rising systems complexity and unabated data growth.

This is the reason why EuroCloud has decided to publish a new Whitepaper to cover this very specific and highly technological matter. EuroCloud's focus is always on supporting the development of a healthy balance between the Cloud Provider and Consumer. This is why we focus on topics such as, legal compliance and barrier free cloud as well as more technical matters.

We would like to support users and suppliers of cloud systems to understand the aspects and tricky areas of the measurement of Application Performance. Hence, we have created this guide, which has been translated into a number of languages.

EuroCloud would like to extend special thanks to Mr Reinhard Travniček, Managing Director at X-tech, and the author of this guide. Mr Travniček is an expert in the field of virtual desktop environments and has been able to draw upon his personal experience and considerable knowledge to produce this insightful guide.

Vienna, September 2014

Dr. Tobias Höllwarth
Vice President EuroCloud Europe

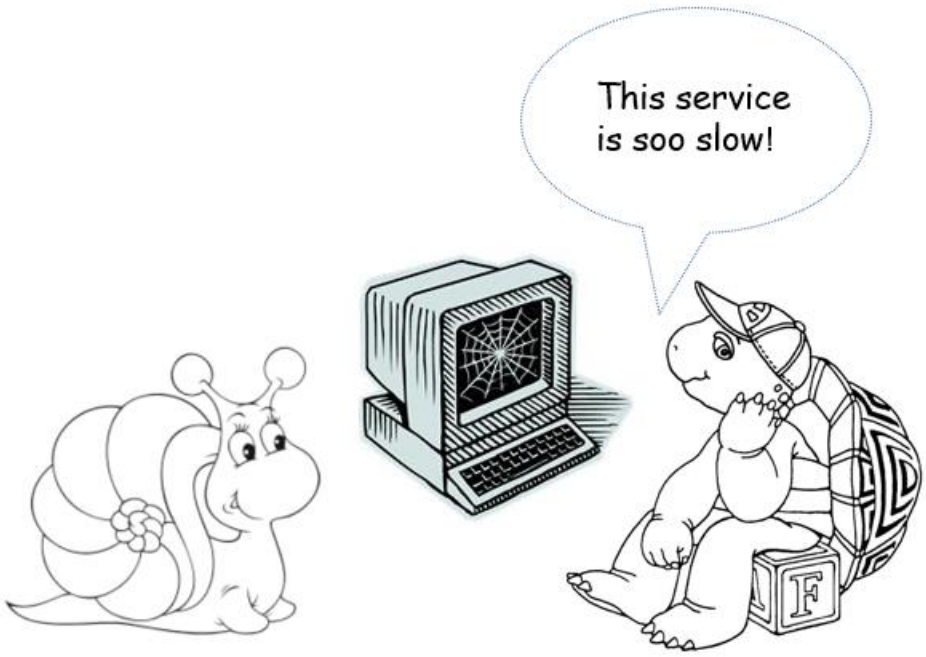
1 Management Summary

Numerous surveys have shown high availability of applications as the top priority for business users, customers and CIOs. However, the more business processes come to depend on multiple applications and the underlying infrastructure, the more susceptible they are to performance degradation. Organisations are increasingly realising the importance of monitoring and improving performance of business-critical applications. Their effectiveness in executing these strategies heavily depends on the approach that they are taking when evaluating application performance and end-user experience. Whilst end-user experience is hard to quantify, application performance data could be collected more easily. Translating these measurements into KPIs and into an end-user experience based SLA is difficult and almost impossible with today's monitoring solutions.

This paper tries to explain the various scenarios where detailed network- and process-information could aid in overcoming the above highlighted SLA shortcomings. And where such tools help in performance troubleshooting, performance diagnoses, root-cause diagnoses and the quantification of end-user experience in VDI / DaaS environments. Nothing shines a light on an IT team's success or failure as much as the application performance and availability does. With uptime as their priority, they need to adapt a more holistic approach to performance management and decision analytics. Through best practices, they can help their companies leverage IT investments to discover, interpret and respond to the myriad events that impact their operations, security, compliance and competitiveness. IT cannot depend on the help desk to be its eyes and ears for keeping systems running at their best.

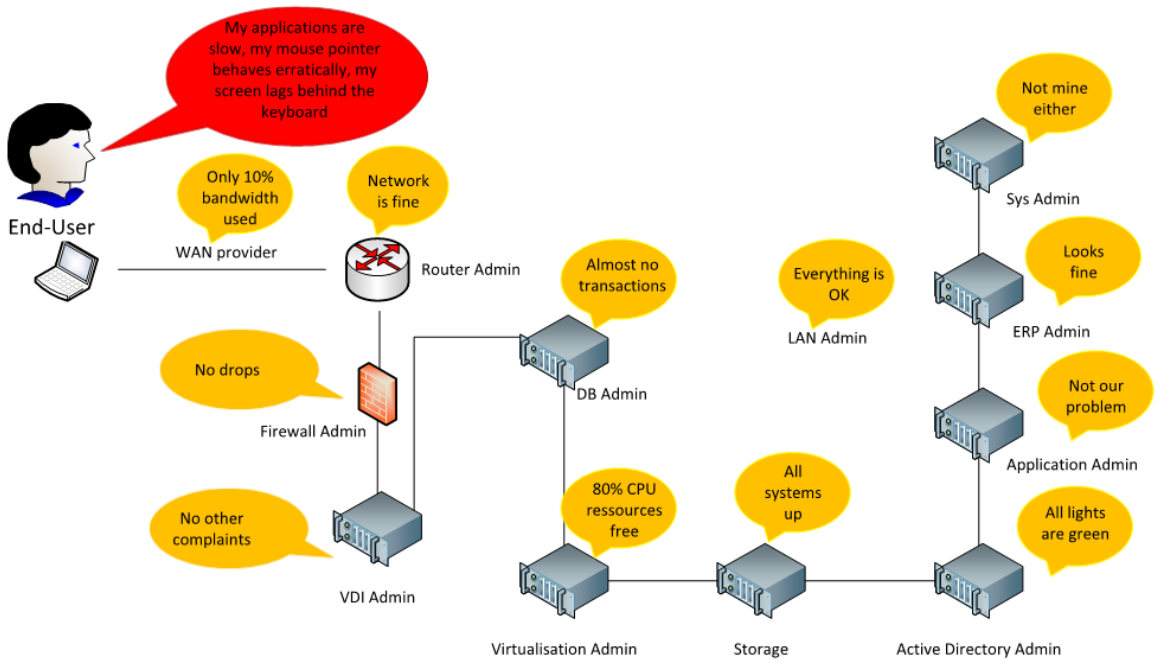
2 Application Performance Management (APM) Today

One can understand that a modern IT environment includes an enormous number of components and clearly a lot of network connections. The combined latencies of these components and networks most likely have the strongest influence on the end-user's "experience".



One of the primary reasons why end users tend to notice application performance degradation before the IT organisation does, is because IT organisations have an ocean of management data to analyse. This data comes from a continually growing number of business transactions as well as from a complex, distributed IT infrastructure that is both physical and virtual, and which is increasingly provided by third-party cloud computing vendors.

Due to the many components involved an end-user complaint will likely be handled like this:



IT groups still acquire discrete tools that focus on a particular hardware, network or software issue, many organisations have ended up with a patchwork quilt of point solutions that do not work well together. While each tool might be indicating that performance of a particular segment or component is “normal”, outages persist and the actual user experience continues to disappoint.

The above illustration shows that solution delivery managers (or even worse the end-user) are not in the position to deal with a large number of external and internal “service providers”. They are simply not capable to understand the various relationships and interdependencies. Consumers of IT services are only interested in what the service can do for them, not in how it is created or managed.

3 The future of APM

3.1 Near Term - Combining all the monitoring methods

A combination of active and passive monitoring (chapter 6.4 has more details on monitoring methods) will give the operators (service providers) a better and complete insight into the state of their system and a way to define KPIs and SLAs to their customers. The active monitoring component will deliver the exact values for KPIs such as, service availability, application startup time, etc. and the passive monitoring part will deliver statistical data for end-user experience prediction and will aid with root cause analysis of the end-user's performance issues? Mid Term – Enhanced passive Monitoring and Big Data Analysis

The proliferation of server virtualisation and the tools needed to monitor and manage virtualised dynamic infrastructures and highly distributed application architectures only expand the data points and metrics that need to be analysed.

Consequently, vital information is often overlooked, resulting in missed opportunities to uncover hidden patterns, relationships and dependencies. Additionally, whatever data is gathered is not normalised or time synchronised, making analysis and rapid problem resolution impossible. Yet pouring more data into obsolete analytics tools only compounds the problem, so a new approach is needed.

The challenge of having oceans of data to analyse is common and has a new name: Big Data. Big Data refers to having data sets that are so large and complex that it is difficult, if not impossible, to process and analyse them using traditional techniques. Companies are developing a wide variety of techniques to analyse huge data sets. In a business environment this is commonly referred to as Big Data Analytics. When it comes to application performance management in an IT environment, this is referred to as Application Behaviour Analytics.

An Application Behaviour Analytics application will automatically learn from live data what normal behaviour is and implement sophisticated statistical techniques to identify situations that deviate from normal behaviour and to identify the cause of that deviation. Implementing Application Behaviour Analytics will enable IT organisations to get improve in both proactive and reactive application performance management. Computer-based analytics will enable IT organisations

to efficiently and effectively analyse this data and it will allow business managers to have the confidence that their critical processes will perform as expected.

We expect to see proprietary APM (Application Performance Management) products to use Big Data methods to enter the market by 2015.

3.2 Long Term – Open Performance Analytics Framework

We believe that there are initiatives that intend to offer an APM platform (as an open source project) by late 2016. The initiative needs to invite all other vendors to use this open architecture to deposit their specific data so that customers can perform Big Data Analytics and obtain business operational intelligence by harvesting such deposited data.

Goals are:

- Adding value to the vendor's product by redirecting development time and effort spent collecting commodity data
- Every vendor has access to every other vendor's data allowing for cross-domain root cause analysis and mashups
- Each vendor becomes part of a much larger ecosystem that is a credible replacement for legacy management frameworks
- The data store will be open source and permissively licensed, protecting vendors from lock in or other adverse consequences

The discussed framework will represent a more holistic approach to adaptive, proactive and predictive operational data management and analysis. The framework should link advanced performance management and Big Data analytics technologies to enable organisations to gain deep and real-time visibility into, and predictive intelligence from, increasingly complex virtualised and mobile systems across the entire application delivery chain.

These new performance analytics techniques must help the enterprise in three ways:

- First, enterprises need to understand what levels of performance (i.e. speed and availability) are needed from their increasingly cloud-based and mobile applications in order to deliver fast, reliable and highly satisfying end-user experiences.

To better understand the properties of the components and their place in the overall application delivery chain, a higher-level assessment of the relationships to each other is required, as well as to the entire system and environment. A comprehensive performance analytics platform will provide visibility across the entire application delivery chain – from behind the firewall and out to the Web, including third-party cloud providers.

- Second, the "point of delivery", which is where the user accesses a composite application, is the only perspective from which user experience should be addressed. As such, the most relevant metric for any IT organisation is not about infrastructure utilisation. Instead, it is at what point of utilisation the user experience begins to degrade. Enterprises need to measure the true experiences of their most important end-user segments, including those that are remote and mobile.

- Third, to provide insights that line of business users can understand and value, IT must establish an effective link between performance management and analytics.

4 Today's and future VDI Cloud monitoring solutions

Monitoring technologies typically fall into two distinct categories

- Solutions for active monitoring, also known as "synthetic" monitoring
- Solutions for passive monitoring, also known as "real-user experience" monitoring

Both of these methodologies have their advantages and help organisations to measure application performance. Active Monitoring tools combined with a system to gather system metrics are used to build the majority of all systems deployed today.

4.1 Active Monitoring

Technologies for active end-user experience monitoring are based on generating simulated transactions to measure application and transaction response times. Basically what happens is, a monitoring agent sends periodical requests across the network, the response time

of the application servers and network infrastructure are then recorded and evaluated.

This synthetic monitoring approach is very effective in addressing the primary goal for application performance management: The prevention of performance issues before end-users are impacted. Organisations therefore proactively test parts of their infrastructure to identify potential performance bottlenecks before they cause disruption of business processes.

Synthetic monitoring gives almost no insight to the end-user experience if used in VDI / DaaS environments, currently a common cloud provided service.

Compared with a classic client / server environment, VDI /DaaS environments include the so called Presentation Layer, where the graphical representation of the software running on a worker is transferred as a “picture” to the client device. This process is very similar to the way TV broadcasting works.

Active Monitoring solutions are mostly based on the recognition of windows events (creation / pop-up, close) and therefore must be completely redesigned for VDI/DaaS to a “pattern recognition” approach.

Since “pattern recognition” is much more time consuming and a lot more prone to error, due to window overlaps, this method is very unreliable if used for VDI / DaaS. That said it is also very likely that the monitoring agent could only be installed on a limited set of operating systems. This will then not allow the prediction of how the performance and user experience, on mobile devices or thin clients will be

4.1.1 Benefits of active monitoring in VDI / DaaS an environment

The real benefits of active monitoring are:

- There are exact measurement points
- The measurement delivers an exact numerical value
- The large number of repetitive measurements helps in trending
- KPIs can be easily quantified
- Find out the performance impacts of tuning parameters

- See the impact of changes or updates such as service packs or security fixes

4.1.2 Drawbacks of active monitoring in a VDI / DaaS environment

There are several drawbacks which include:

- Agent consumes more CPU resources because of pattern recognition
- Agent cannot install on the worker in the presentation layer because there are no patterns to recognise
- If the software does not respond with the exact “predicted” pattern, the measurement will fail
- The agent most likely uses a synthetic user
- The measurement setup is not reflecting real life, because the synthetic user uses a very static environment to perform the measurement
- Active Monitoring will not cover end-user experience because it does not cover:
- Degraded audio performance
- Asynchronous audio / video
- Erratic mouse movement
- Lagging keyboard
- Reduced video frame rate
- Scrolling problems
- There is almost no way to deduce the root-cause of performance degradation

Today’s active monitoring tools are often combined with tools that report process-level metrics, such as CPU, memory, and IO. Even though this metric allows IT departments to measure availability of business services that are being delivered to end-users, it does not allow them to measure the actual quality of IT services that business users are experiencing.

Monitoring application uptime is no longer enough to ensure seamless delivery of business services. Active monitoring tools do not deliver on end-user experience in DaaS environments.

4.2 In Future: Passive Monitoring

A passive monitoring solution installed will provide IT Ops with application- and virtual desktop performance correlated with user productivity. Armed with this data, IT Ops can rapidly investigate users' complaints of poor app performance, determine other impacted users and the likely root causes resolving the issue before workforce productivity is impacted. For that reason, organisations will increasing using performance metrics such as application and transaction response times and end-user experience indexes that allow them to monitor the speed of applications and to evaluate the quality of the end-user experience. Plenty of CPU resources, network bandwidth to spare, affordable storage space and Big Data analyses will help to build passive monitoring systems and solutions which are able to deal with the huge amount of data gathered by the monitoring sensors installed.

4.2.1 Functional requirements

The monitoring solution needs to meet the following criteria:

- Complete Traffic Visibility: All transactions between applications /services and each infrastructure element must be monitored. The solution must be agnostic to applications, coding, OS, virtualisation and infrastructure.
- Stack Visibility: to build the topology map and to follow all transactions through the stack from end-to-end full stack visibility from layer 2 to layer 7, networking and storage is needed.
- End-to-End Visibility: agents on both servers and clients, SNMP and WMI monitors allow to view application /service performance as seen by the end user. The entire dataset is sent back to the master for event correlation, root cause analysis and resolution proposal.
- Topology Discovery: As the applications and services start to interact with each other and build a workflow, the monitoring solution should build the topology and understand the end-to-end and hop-by-hop topology of the virtual and physical mappings.
- Root Cause Analysis: Based on many thousands of stored instances and on rules, which are continually being incremented by auto learning, the root cause analysis of an event could be deterministically established.

4.2.2 The architecture

The monitoring system will most likely have a master system containing a database and a management GUI. There must be granular alerting and an extensive reporting tool. For alerting it is necessary to establish baselines. Modern systems try to assist the administrators with auto-baseline. Additional trending software would be appreciated. Data will be gathered by agents or sensors, which could be installed on the monitored component itself, or using remote sensoring capabilities like SNMP or WMI queries while running on the master system.

The entire communication between the agents and the master systems must be encrypted, and secured. The communication path should use standard ports (443) and standard protocols (https), to allow the agent to traverse proxies and firewall secured environments.

4.2.3 Where to place agents and which data must be gathered Sensor on the VDI client

VDI client devices running complex multi-tasking operating systems, could introduce performance and user experience problems of their own. Simpler client operating systems like IOS or Android with only one foreground process are less prone to self-inflicted problems. On the other hand complex operating systems give room for more sophisticated monitoring agents.

There is no doubt that a monitoring agent, at least recording network connections, latencies and transactional behaviour, on the VDI client is needed. On multi-tasking operating systems the monitoring agent could aid troubleshooting by recording and reporting the system status (CPU, Memory, Disk IO and running processes)

- Measuring VDI client software startup

Although the client software startup time gives no important insight into the performance of the system, this process is the first visible detail to the end-user, so if possible the VDI client startup time, which normally includes an authentication and a presentation phase should be measured. Once the remote desktop / application is selected by the consumer, session and client context initialisation takes place, so measurement must also start on the presentation layer.

- VDI client context initialisation

If the session is initiated there is some handling of client devices (Printer, drives, etc.) this might cause considerable delay. The monitoring agent should be able to record those delays and report back to the master.

Sensor on the worker

If desktop / application from a VDI/DaaS infrastructure is started, a session on a worker is instanced. The worker will carry out the execution of the programme (s) available on the desktop or called from the remote application definition. Therefore an agent on the worker is needed.

- Measuring the Session initialisation

To set up a session (context) a complete “login” process needs to be performed. During the login process, the processing of Group Policies, Logon Scripts, the connection of various redirected devices (printer, local drives, USB, etc.) influence the desktop/application startup experience of the end-user. From what is said, what do we need to know to understand the end-user experience?

We have to understand what the logon process is doing, what is causing the session setup process to be slow:

- there is some traffic to the Active Directory (RFS),
- there is Fileserver CIFS (CFS) traffic for the logon scripts and Group Policies,
- there is a lot of CIFS traffic to the roaming profile share,
- there is some traffic to the users base directory (Homeshare),
- there is time needed to process the Group Policies,
- there is time needed to process the logon script and
- there is, of course, CPU time needed to “build” the session context on the worker.

Technically this means there is only one connection between the VDI client and the worker, but many network connections from the worker to the backend services (both CFS and RFS). Therefore we need a monitoring agent /sensor on the worker and if possible on all backend systems.

Since the session on the worker could be attributed to an end-user, it could aid to the overall picture if the monitoring agent records this information with the master. Most likely the end-users will not complain

about the “published desktop” being slow, they will be more precise, for instance: “my ERP system is not working”, or “my videos from the web are slow”, or “my Office application is slow while loading data”.

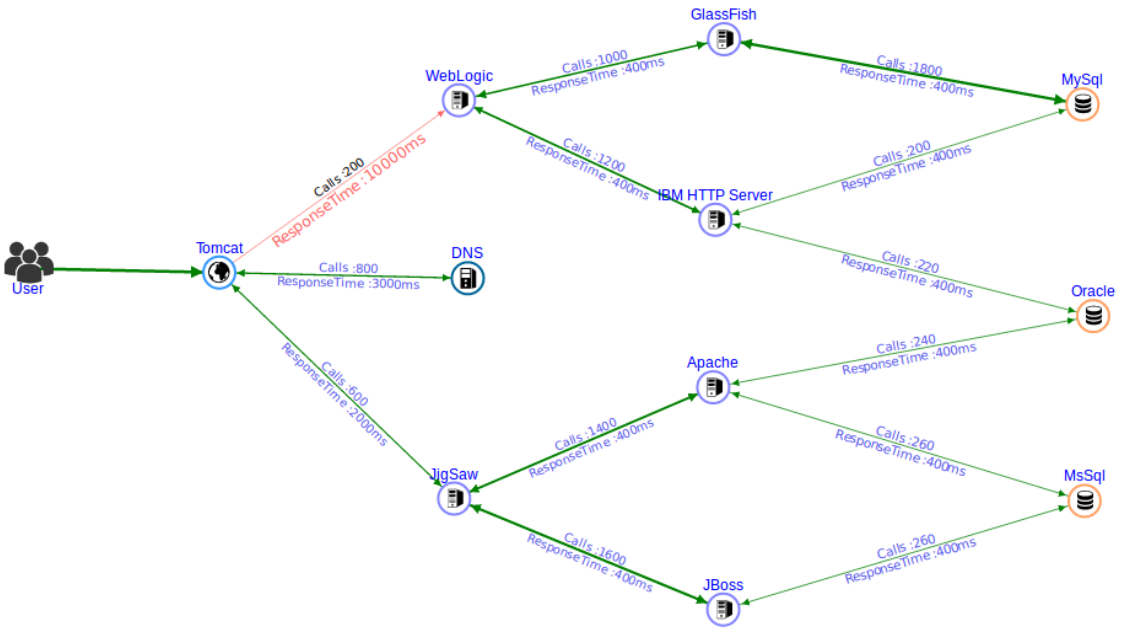
- Application performance

Performance of applications running within a session context on a worker is dependent on several factors:

- CPU cycles on the worker
- Available RAM in the worker, because paging is causing delays
- Network bandwidth from the worker to the backend infrastructure
- Transactional delays with backend services (both CFS and RFS)
- Performance bottlenecks on the hypervisor
- Performance bottlenecks on the storage systems attached to the hypervisor

So once again, to monitor the performance of a single application you must correlate data from a variety of monitoring agents / sensors. At that point Big Data analyses is key. The topology of the application must be understood, while this knowledge (dependencies) could be entered into the monitoring system by hand, good systems will perform an automatic topology discovery.

Figure 1 - Application Topology



Sensor on backend systems

From the topology view above one can see that it is important to install monitoring agents on the backend systems themselves. Many customer facing services (CFS) like file services are depending on resource facing services (AD for authorisation, DNS, etc.) and therefore a root cause analyses is only possible if all the transactional data of such services are recorded and stored in the master database.

4.2.4 Additional Sensor/Agent function - inform the end-user

The author believes that passing status information to the end-users is crucial, therefore the implemented solution must provide feedback on the system status to the consumers.

Proactively informing the end-user is a way to reduce service desk calls. In the event of system degradation the monitoring agents GUI on the client device should be used to notify the end-user. For the network connection this could be achieved by showing the latency as a number and a colour indicator (red, amber, green) and for backend services this could be achieved by a balloon pop-up which provides the status of resolution measures and timelines so users do not have to call the service desk for updates and can plan their workload accordingly.

4.2.5 Benefits of passive monitoring in an VDI / DaaS environment

There are several benefits on passive monitoring:

- Data from real life users
- Data from every user
- Data from a large number of agents/sensors
- Transactional performance data to predict end-user experience
- Easier root-cause analyses
- Trending for all the systems involved

4.2.6 Drawbacks of active monitoring in an VDI / DaaS environment

The drawbacks are:

- No measurement like “Logon Time”

- No repetitive measurement of always the same application interaction
- Possibly no immediate alert if an application fails

5 Benchmarking, Load- / Stress-Testing and Regression Test

Currently, benchmarking, load testing and regression tests are most likely performed with active monitoring systems. Load and stress testing answers the following questions simply and conclusively:

- Will my application fall apart under load?
- Will my pre-defined performance targets be met?
- How many simultaneous users can my website handle?

As already mentioned in the drawbacks of active monitoring systems, the user context and the simulated unique user workloads are synthetic, the data gathered are only good for comparison with other systems and or for trending. The answer to the question of when the system will become overloaded might be possible, but, the test will not point to the root cause of the overload.

A regression test could be performed to find out the impact of:

- newly introduced tuning parameters
- changes or updates such as, service packs or security fixes
- changes to the architecture
- changes in CFS or RFS
- hardware upgrades

6 More detailed view based on Desktop as a Service (DaaS)

New distributed computing architectures and approaches to agile application development have made computing far more scalable and dynamic than ever before. Even with all components (except the client device) on-premises VDI /DaaS environments are considered to be very complex.

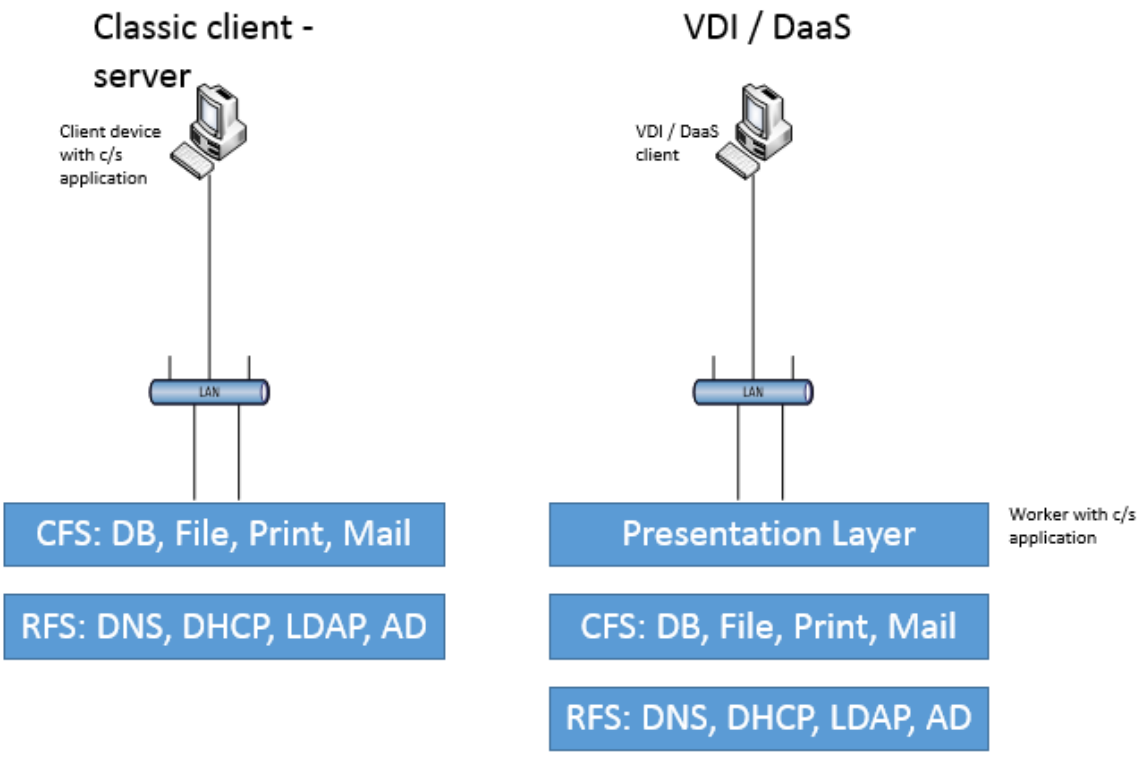


Figure 2 - Client Server versus VDI

Compared with a classic client server environment VDI /DaaS adds one additional layer, the so called Presentation Layer. Unfortunately for the responsible management team, typical VDI / DaaS solutions are targeted to WAN / Corporate WAN environments which add additional components into the communication path between the client and the presentation layer.

Discussion with VDI / DaaS consumers are showing, the following typical SLA shortcomings:

- Failure to adequately define the service(s) being offered
- Failure to define how a service is measured
- Exclusions for scheduled maintenance without defining maintenance frequency, notification, time of maintenance window
- Failure to distinguish between partial failure (such as slowdowns, reduction in capacity, inability to launch new

desktop sessions) and total loss of service and with it, all unsaved work

- Exclusion of equipment dedicated to an individual customer over which the customer had no direct control
- Exclusive assignment of service monitoring responsibility with customer
- Failure to account for service provider induced error
- Failure to account for infrastructure software failure
- Failure to account for data centre connectivity failure
- Failure to document compensation levels for failure to meet service level agreements

The number and breadth of SLA weaknesses suggests that a substantial number of VDI/DaaS providers (both internal and external) have some way to go before they fully understand the nature of the service that they offer.

DaaS must be the same as any other utility; just as a user does not need to know the complexity of the power grid to understand how to use electricity to heat or light their home, so a DaaS user should not have to understand the technology that is used to provide the service.

The position of the client device (WAN / LAN) and the usage of virtualisation (servers / network / storage) and the integration of hybrid clouds defines the complexity of the environments. Nevertheless, complaints raised by end-users (mostly about user experience issues) are identical.

What might a user complain about?

- Slow startup time of VDI client component
- Slow application startup
- Sluggish performance of the application
- Erratic mouse and keyboard
- Slow / imprecise scrolling in applications like Microsoft Excel
- Asynchronous video /audio
- Reduced Video Rates
- Scrambled audio

6.1 VDI / DaaS with LAN based client (simplest)

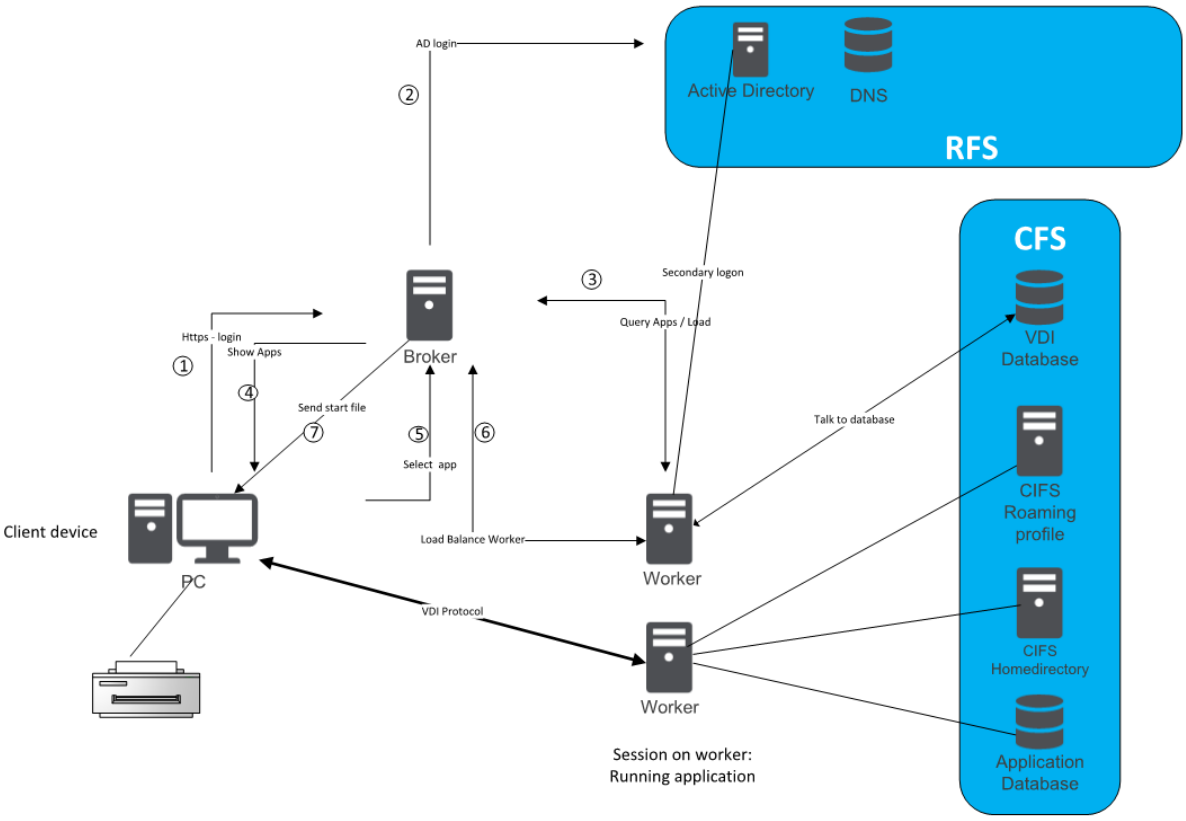


Figure 3 - VDI / DaaS LAN based

In this pure LAN based environment, the entire communication between the client device and the worker could be seen; there is no hypervisor or storage subsystem involved.

6.1.1 What is causing performance problems in this scenario?

- If login is slow, then it is most likely that the Broker server itself is overloaded and therefore the root of the problem. The interaction with both AD and the XML query(s) are light weight, but nevertheless an AD or XML outage of a configured server could lead to a significant (30 seconds) timeout delay during VDI client startup.

- Application startup involves a complete login process (at least for the first application started on that worker and for that user). The login process itself is complex and puts a burden on the worker. The network interactions with the CIFS share holding the Roaming profile are high.

- If the application does not perform as expected, it could be a resource problem on the worker (CPU, RAM) or a network or process latency in the backend. RAM could cause a problem, because of paging, once the worker uses its hard disk for paging the application experience for the end-user starts to vary.

- Complaints about mouse and keyboard responsiveness should never happen in LAN based deployments, because these are very small data packets and the connection is in the Gigabit/s range.

- Slow / imprecise scrolling is also mostly seen in WAN scenarios, but could also happen if an extremely under-powered under- powered client device is used.

6.2 VDI / DaaS with LAN based client (virtualised)

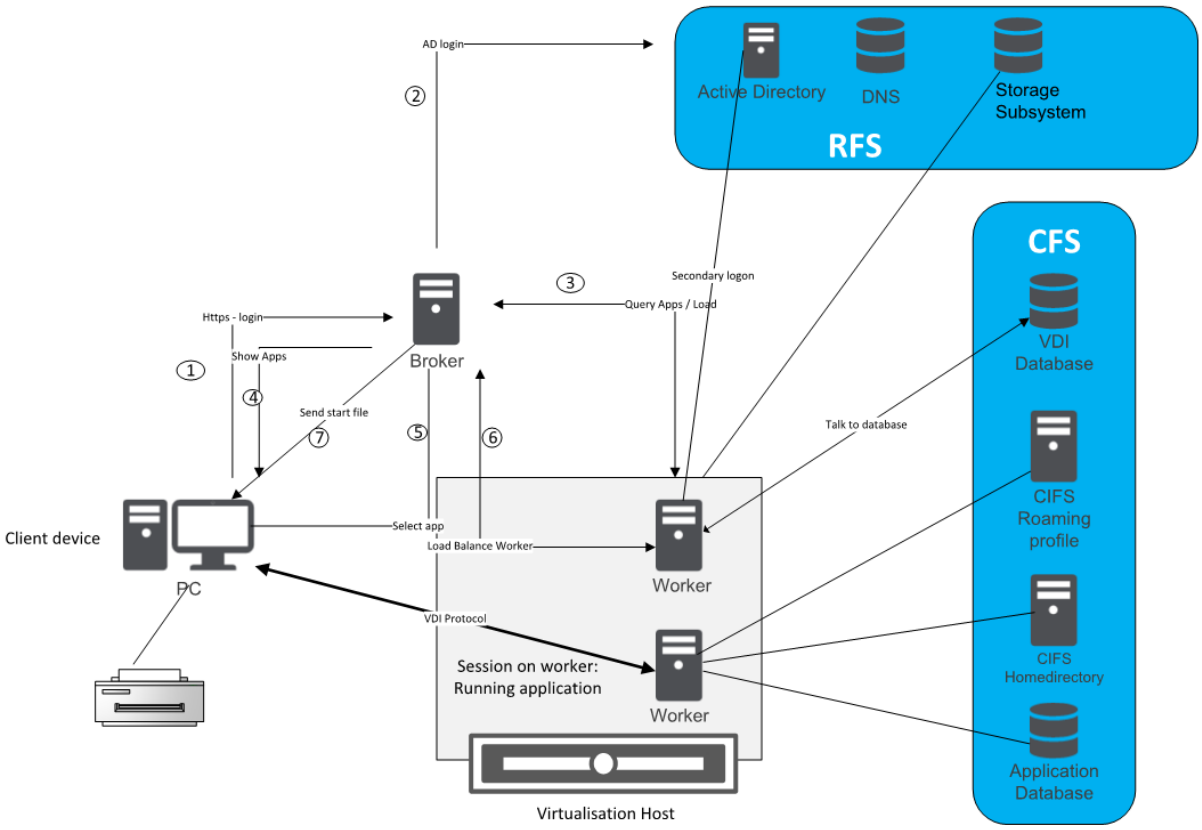


Figure 4 - VDI / DaaS LAN based with virtualisation

In a pure LAN based environment, the entire communication between the client device and the worker could be seen, but there is now a virtualisation layer (hypervisor) and a storage subsystem/network involved preventing this.

6.2.1 What is causing additional performance problems in this scenario?

Reasons to complain are identical to the first described scenario. The monitoring complexity has grown, because of the hypervisor and the storage system involved. With hypervisor and storage in place there are more systems which could introduce latency to the system. Even if the end-to-end transaction time of a particular app is excellent, if the underlying platform is sluggish – perhaps due to CPU power, memory availability, or other background processes that are resource hogs -- the end user's experience, with what should be a well-performing application, will be poor.

6.3 VDI / DaaS with WAN based client (more complex)

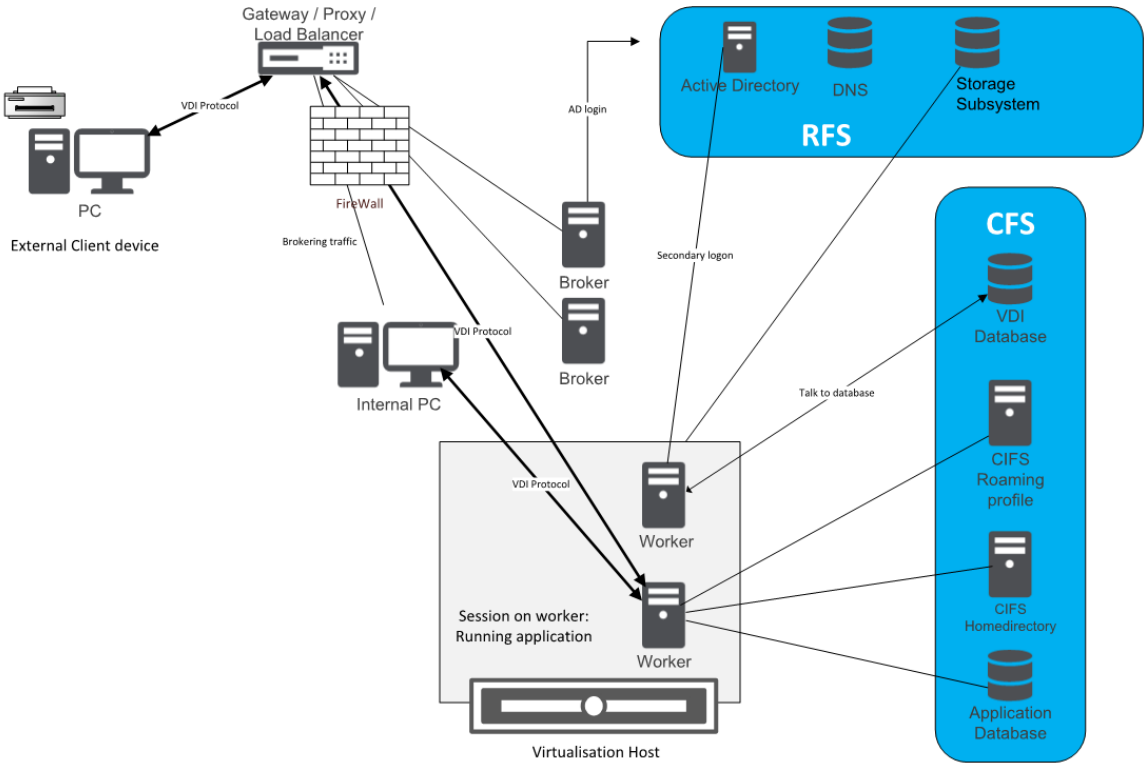


Figure 5 - VDI / DaaS WAN based client

In this scenario there is a remote access a Gateway /Proxy, so the entire communication between the external client device and the worker could *not* be seen, there is a hypervisor with associated storage, and the Broker is redundant with Load balancing on the Gateway in place. Typically the Gateway / Proxy is located in a DMZ, so all the traffic has to traverse the firewall system as well.

6.3.1 What is causing additional performance problems in this scenario?

The reasons to complain are identical, but the communication between the external client device and the worker has to traverse two firewalls and is proxied by the Gateway. The Gateway itself is used for authentication and authorisation so there is an additional need to monitor the network / performance from and to AD.

If users start to complain about mouse and keyboard responsiveness the root cause will be network latency between the client device and the gateway. Slow / imprecise scrolling will be seen if there is no sufficient network bandwidth available between the client device and the access gateway.

6.4 Terms and Definitions

As some of the terms are used differently in other IT contexts, it is necessary to define their meanings to enable a greater understanding of this Whitepaper in this context the following apply:

- Client Device

Client device is the end users device, which is used to access resources published by a VDI / DaaS provider.

- VDI Client Software

VDI client software is the component responsible for communication from the client device to the VDI / DaaS infrastructure.

- Client Resources

Client resources are printers, hard disks, CD-ROM, audio, USB, and so on. All of these resources are local to the client device.

- Gateway

Gateway is a secure proxy for the data traversing the WAN from and to a remote client device. The gateway most likely will perform authentication and authorisation of the end-user.

- Worker

A worker is either a remote session host (aka terminal server) offering a session, or a Windows Client OS controlled by the VDI / DaaS software.

- Session

The context on the worker. Within the session an application is started. A session (context) exists after the complete "login" process to this particular worker has taken place, and depending on the system, the session either shares the worker with somebody else (Remote Desktop Host) or the worker is dedicated to the session (Workstation OS).

- Web GUI

Web GUI provides the end-user with a graphical presentation of the VDI / DaaS resource available.

- Broker

The Broker is providing information about available resources, least loaded workers, etc. The broker in most VDI / DaaS environments is responsible for load balancing and reconnecting to existing or disconnected sessions.

- CFS[1]

Customer Facing Services are higher level services, such as, File-, Print-, Database-Services, CRM application host, etc.

- RFS[2]

Resource Facing Services are backend processes, such as Active Directory, DHCP, DNS, LDAP, etc.

- Physical or virtual infrastructure components

Infrastructure components are the servers and hypervisors running the worker, storage subsystems, networks, routers, firewalls, etc.

[1] Customer Facing Service (NGOSS): http://tinyurl.com/k65xzn2

[2] Resource Facing Service (NGOSS): http://preview.tinyurl.com/k65xzn2

7 Legal Notice

7.1 General Information

The information provided in this guide serves the purpose of illustrating special legal aspects related to cloud computing in a general manner. It does not represent legal advice and is not intended to replace such legal advice, as this always requires precise knowledge of the individual factors and, in particular, the concrete case at hand.

7.2 Contents of the Guide

The publisher/the authors do not accept responsibility for the completeness, accuracy, or currency of the information provided. This specifically applies to new developments in case law or the legal framework. Claims for damages against the publisher/the authors related to material or non-material damages caused through the use or non-use of the offered information or through the use of incorrect and incomplete information are generally excluded.

7.3 References and Links

As regards direct or indirect references to external content (for example, 'links') outside the sphere of responsibility of the publisher/the authors, liability would only apply if the publisher/the authors were aware of the content and if it had been technically possible and feasible to prevent the use of illegal content. The publisher/the authors hereby expressly declare that at the time when the links were placed, no illegal content was detected on the linked websites. The publisher/the authors have no influence on the current or future design, content, or ownership of the linked pages. They expressly disassociate themselves from any content of all linked pages that were modified after the links were placed. Any liability for illegal, erroneous, or incomplete content and any damages resulting from the use or non-use of this information offering shall rest solely with the provider of the referred website and not with the party merely providing links to the respective publication.

7.4 Copyright

The contents, texts, graphics, and images included in this guide are copyright protected under Austrian copyright law. Any utilisation that is not permitted according to copyright law shall require the previous written authorisation of the publisher. Contributions from third parties are marked as such. This applies in particular to the reproduction, editing, processing, or rendering of content in databases or other electronic media. The unauthorised reproduction or disclosure of the guide or any part thereof is explicitly forbidden. Individual or private use is exempt from this; however, private use does not include the right to disclosure to third parties. The same shall apply to publications or other works.

7.5 Fees

This guide is provided to the addressees/recipients free of charge.

8 Author



Reinhard Travniček
Managing Director Technology
X-tech Handelsges.m.b.H.
Hauptstraße 157
2231 Strasshof , Austria

T +43 2287 3928-23
M +43 699 13928 023
Reinhard.Travnicek@x-tech.at

X-tech is a EuroCloud Austria Sponsoring Partner

Photo: Copyright: Michael Fritscher

EuroCloud.Austria

Verein zur Förderung von Cloud Computing

Museumstrasse 5/14

1070 Vienna, Austria

E-mail: info@eurocloud.at

Web: http://www.eurocloud.at