vergic
MAKES IT PERSONAL

Vergic Cloud Computing Platform

Dokument ID: Vergic Cloud Computing Platform

Datum: 20 Sep 2014

Revisionsdatum: 20 Sep 2014, Esben Carlsen

Vergic Cloud Computing Platform

1 Introduction

This document describes Vergic Cloud computing platform architecture Vergic Engage for tenants, shared, dedicated and on premise. The product Vergic Engage is hosted on the Vergic Cloud computing as a SaaS.

Vergic Engage is a highly scalable, redundant and secure SaaS platform. SaaS, scalability, security and redundancy systems are inherently complex and expensive. There are many moving parts that need to work together. The Vergic cloud computing is the workhorse for Vergic Engage.

2 Fundamentals

Vergic Engage has two types of traffic, HTTP(S) and RTMP. HTTP(S) traffic can be split into three groups, namely static content, semi static content and non-static content.

2.1 Cloud computing, IaaS, PaaS and SaaS

Vergic cloud computing offers infrastructure and software as a service as defined the standard defined on Wikipedia. For clarifications on abbreviations and definitions, please look here: short architecture overview of cloud computing is illustrated here:

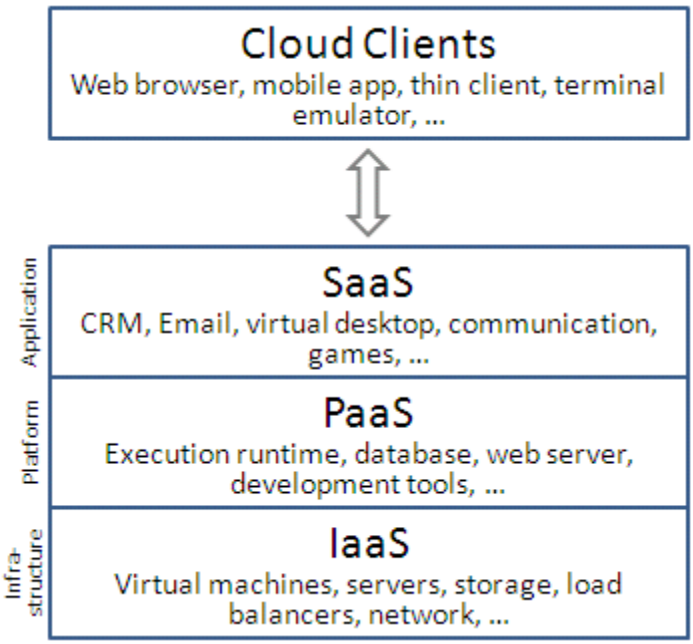

The Vergic computing cloud is build using Microsoft Windows Azure Pack or WAP for short, which is not to be mistaken for Microsoft Azure. More information about WAP can be found here.

2.2 Static Content

Static content is content that never changes, like CSS, JavaScript libraries and images. Static content can be cached for a long time to reduce bandwidth and latency without checking for updates and

without the system breaking. Static content is well suited to be cached client side without the need to check for updates.

2.3 Semi Static Content

Semi static content is almost like static content, but can be changed every now and then. Semi static content can also be cached for a long time, but needs to be checked for updates server side. Vergic Engage does this type of caching using ETag headers in the HTTP(S) protocol.

Vergic Engage produces a generated JavaScript file that is comprised of static content and the current configuration of the system. The configuration is included in the generated JavaScript file to reduce the number of HTTP(S) requests done per browser page load and to reduce bandwidth. Another benefit of this is that configuration changes are almost immediately propagated to the entire system, including clients.

2.4 Non-static Content

Non-static content is REST traffic. REST requests are issued to a web server to perform an operation and return something new every time, caching is therefore not possible. REST traffic in Vergic Engage when executed in a browser is for the most part polling requests and polling requests return change events. Polling requests are asynchronous in nature and therefore even though there is a potential geographical introduced latency overhead the overall user experience is very good, it is actually not noticeable at all.

2.5 SSL Handshaking

Vergic Engage runs secure by default using HTTPS. Using HTTPS / SSL has a processing and time penalty. It is expressed especially in the web world by the frequent connections to a web server because every new connection needs to go through a SSL handshaking process, which is complex and takes time. But by placing an SSL off loader geographically close to the client the SSL handshaking process time can be drastically decreased by simply minimizing the physical data on the wire transfer latency.

2.6 CDN / Cache

All static and semi static content is served via a caching server, so content can be served extremely fast. All content is always stored directly in memory and written directly from memory to a network socket. Having a cache close to the solution enables pushing out changes very quickly. It normally takes less than one second from a change has been made to the configuration, until the change has propagated into the cache.

2.7 RTMP

Vergic Engage supports video/audio directly in the browser without the use of additional plugins. The protocol used is RTMP. If needed a RTMP server cluster can be added in any region as well. RMTP data is considered transient and if video/audio storage is required additional configuration is needed.

3 Hosting overview

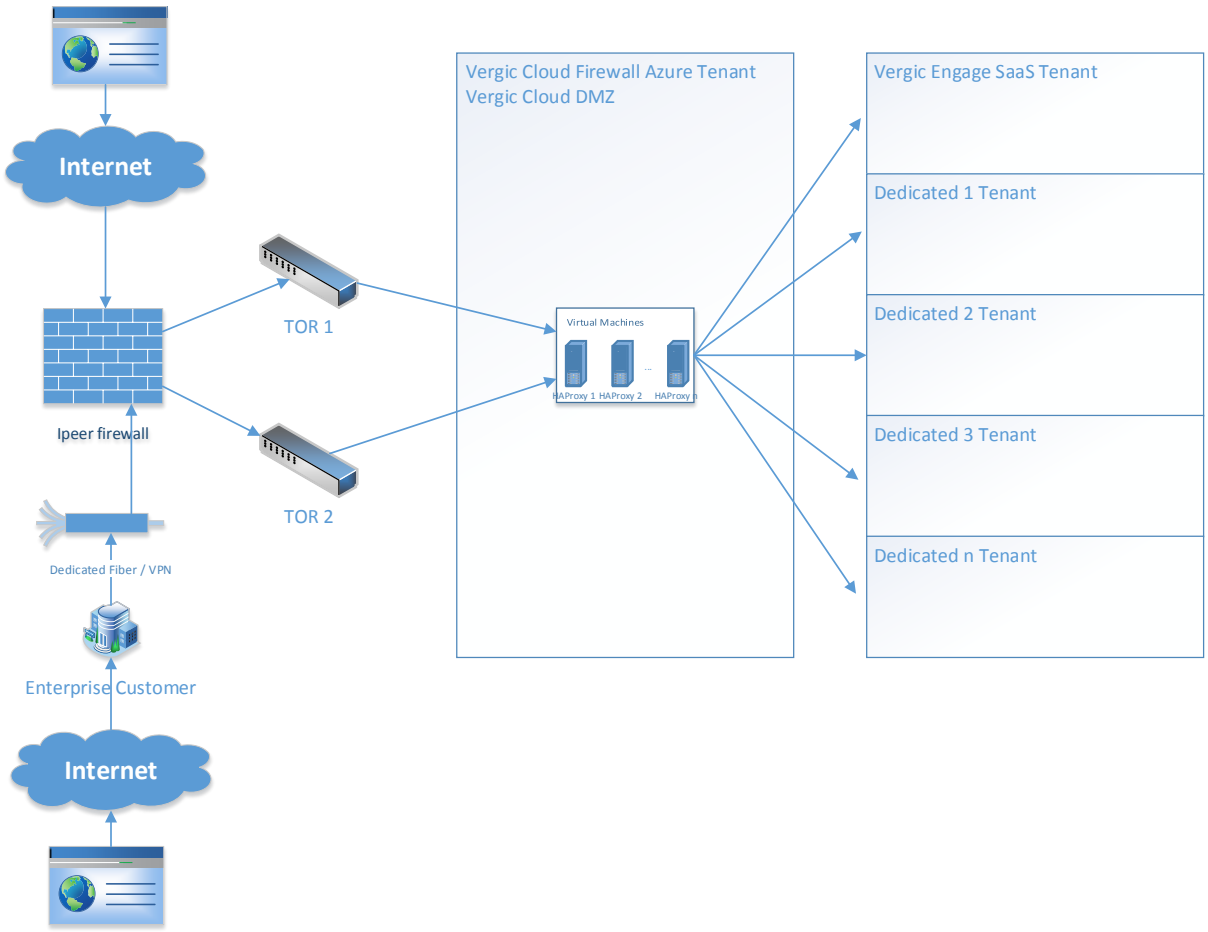

Figure 1

Figure 1 shows a helicopter overview of the Vergic Cloud Computing Platform. Vergic Cloud Computing Platform is a real IaaS. Every Vergic Engage installation is always isolated into a single tenant. Tenants are always redundant because the IaaS is redundant. No Vergic Engage traffic can directly access any network inside a tenant from the Internet, traffic always has to pass through three firewalls, excluding our network supplier Telia that also has (D)DOS protection.

ddos-protection-pro
duktblad-eng-tsp-355

The outer most firewall is our hosting providers firewall, which does layer 4 filtering and does the grunt work on (D)DOS and other commonly known attacks (NSFocus ADS400 Series - http://en.nsfocus.com/2012/ads_pro_0809/26.html).

After passing through the first firewall traffic goes to Vergic Cloud Computing Platform firewall cluster. Figure 2 shows in more detail the setup of the Vergic Cloud Computing Platform firewall cluster.

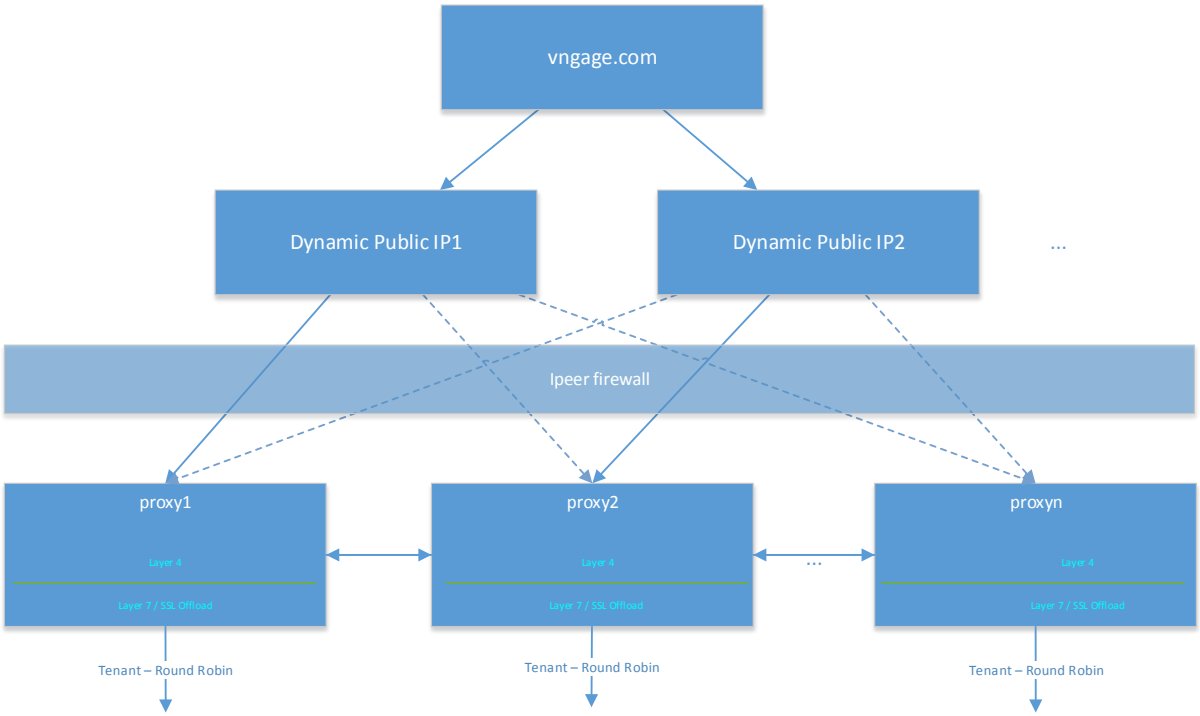

Figure 2

Figure 2 illustrates the configuration of Vergic Engage firewall, load balancer, SSL off-loader and content caching.

At the outer most layer is the DNS. The DNS resolves in at least 2 different public IP addresses. Browsers recognize that when a DNS is resolved into multiple IP addresses they have to do round robin on requests to the resolved DNS hosts. This way Vergic Engage get an evenly distributed traffic pattern on the public facing firewall servers.

If one of the at least two public facing hosts were to crash all traffic will be routed to the other hosts as determined by round robin in the client. This ensures that at least one server can crash at any given time and the other server will still be able to handle the load.

The TTL on the public DNS typically have a low value in the minute region, which enables the ability to quickly modify publically facing host setup.

When a request is received on a host first TCP layer 4 load balancing is performed. Layer 4 load balancing is extremely quick and takes very little time and processing power. Layer 4 traffic load balancing is performed using round robin to all of the available SSL off loaders in the cluster. This allows Vergic Engage to have an infinite number of SSL off loaders with a limited number of publically facing hosts. At any time more SSL off loader can be added to handle additional load.

After SSL offloading is the HTTP traffic. It is determined if the traffic is (semi)static content or REST traffic. If a request is (semi)static the request is routed to one of the cache hosts using weighted round robin.

In this setup any of the servers can crash or be taken down without it affecting the uptime, which makes the solution scalable and redundant.

The dashed lines in Figure 2 represent secondary connections in case SSL off loader primary cache server crashes.

The firewall has the capability to handle a range of attacks like DOS, DDOS, Slowloris and unfair users etc. It can also be configured to handle all or subsets or SSL ciphers as per customer wish. By default we support most major ciphers (SSL3, TLS1, AES, etc.).

After the traffic successfully passes through the Vergic Cloud Computing Platform firewall cluster it is routed to one of the tenants. Figure 3 and Figure 4 shows the redundant setup for round robin load balancing HTTP and RTMP traffic.

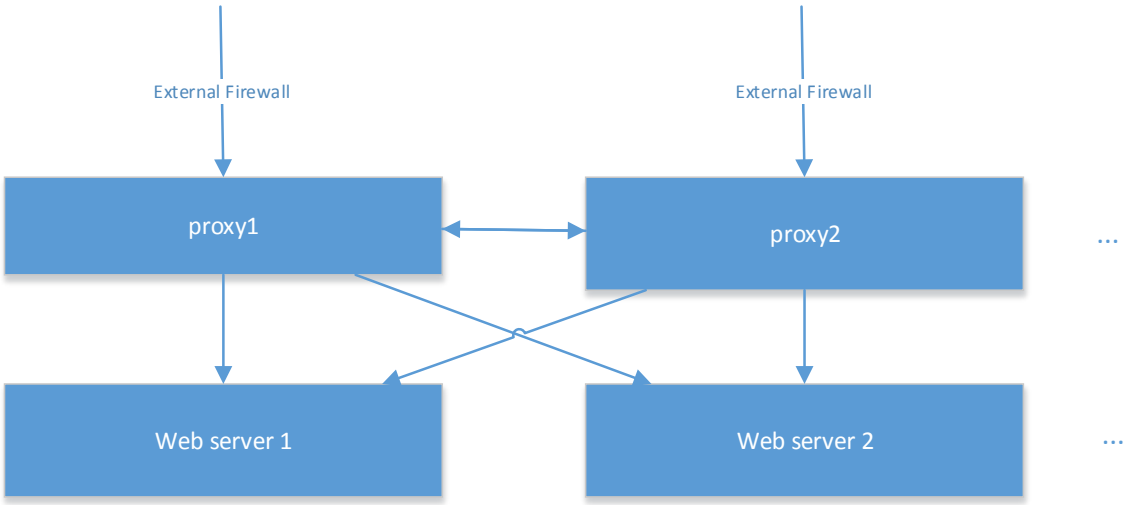

Figure 3 – Tenant firewall / load balancer

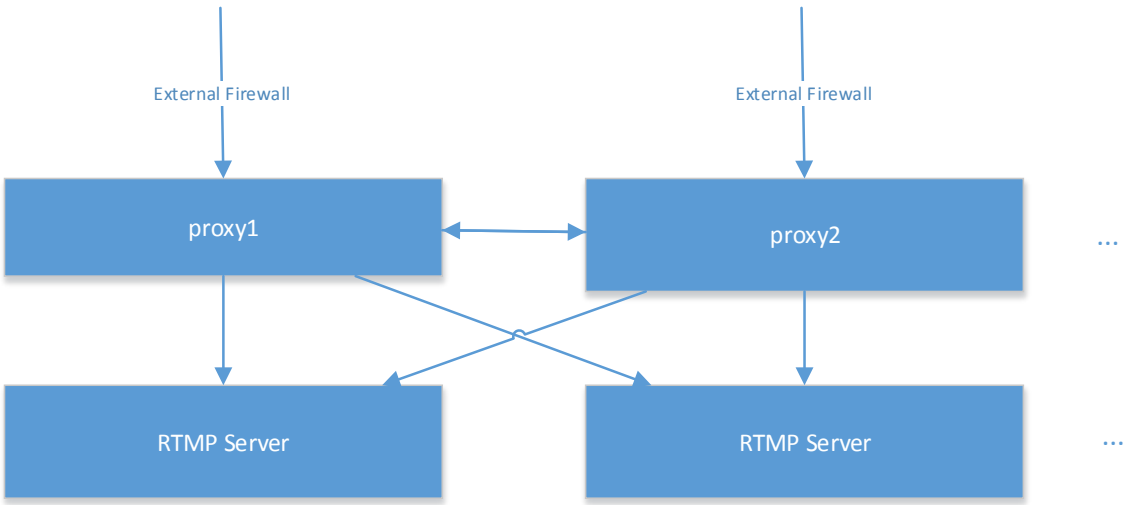

Figure 4 - Tenant firewall / load balancer for RTMP

4 Azure Pack Tenant

All tenants are always configured exactly the same to make maintenance automatic, repeatable and safe. Tenants have three network tiers, each build on the security of the previous.

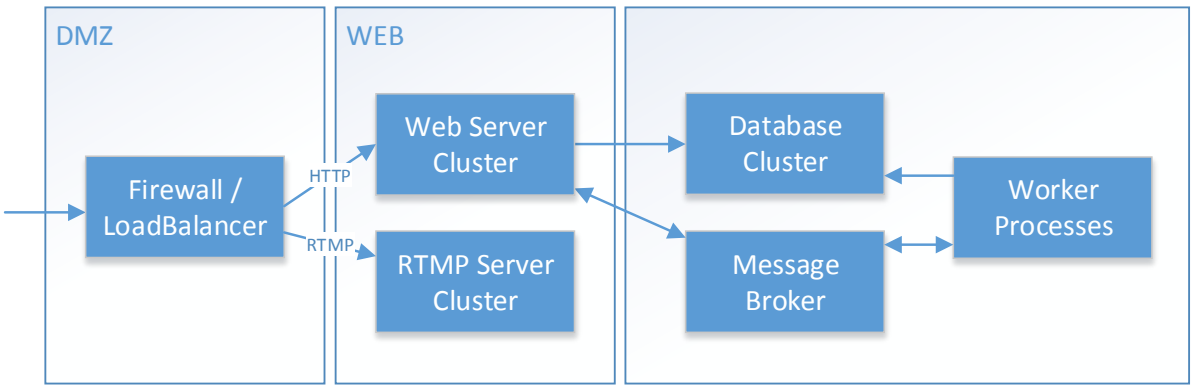

Figure 5

The logical overview of Vergic Engage can be broken down into what is illustrated in Figure 6. The tenants are configured consistently and automatic, and are monitored extensively.

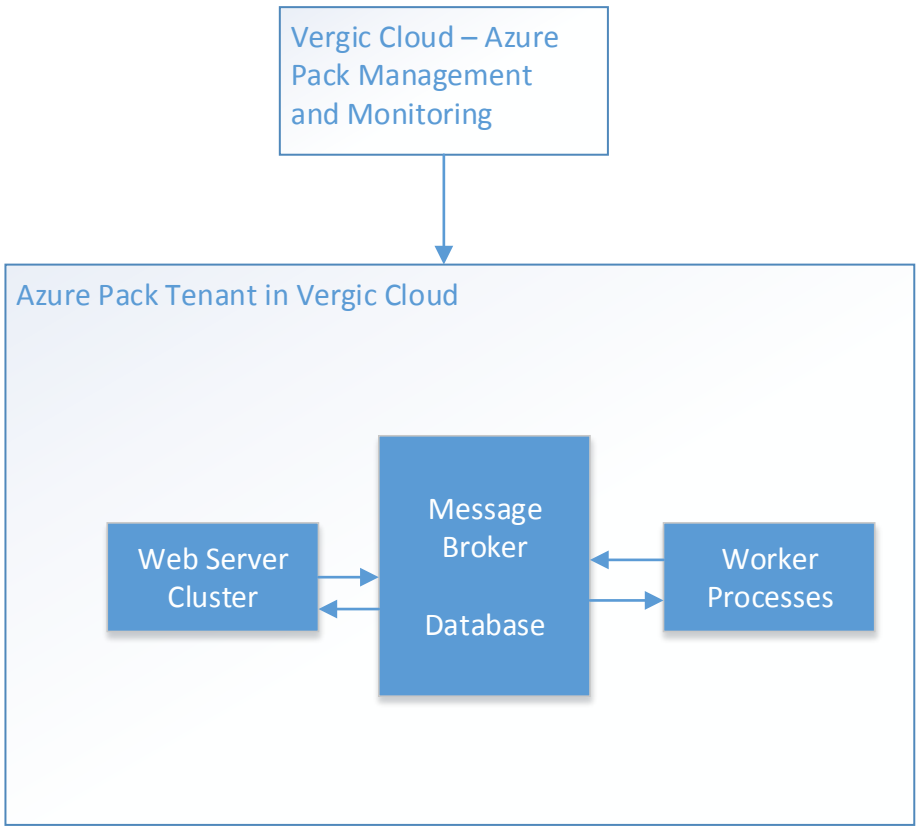

Figure 6

Vergic Engage is in itself designed to be a SaaS product and Vergic Engage can run on a single tenant sharing many concurrent accounts. If preferred Vergic Cloud Computing Platform can create dedicated tenants on shared hardware that exclusively run Vergic Engage for a single account, or even create dedicated tenants on dedicated hardware for single accounts. The hardware can be

placed anywhere in the world and even in multi concurrent co-location sites. The only requirement is that hardware must be part of Vergic Cloud Computing Platform WAP setup.

4.1 Scalability and reliability

In the Vergic Cloud Computing Platform WAP everything is high availability.

The system is continuously monitored for all vital application and server statistics. Technical staff is automatically notified by email and SMS if a threshold is exceeded. Additionally the system automatically can automatically spin up new instances.

Besides high availability everything is redundant, there is at least two of everything. So even if one part of the system fails, it will have no immediate effect.

5 Identity Provider / SAML2

All user authentications towards Vergic Engage are done through the SAML2 standard, which is widely used, and known as a secure system and which documentation is freely available on the Internet.

Any identity provider - that is trusted by Vergic Engage to issue SAML2 security tokens - can be used to authenticate users in Vergic Engage. For example Microsoft Active Directory Federation Services is commonly used as a SAML2 issuing service, usually used as a SSO (Single Sign On) service as well.

Vergic Engage has its own SAML2 issuing service. Vergic Engage identity provider is totally separate from Vergic Engage and is not hosted in the same tenant as Vergic Engage. Any number of dedicated instances of the Vergic Engage identity provider can be installed.

6 Backup

Even though data is always replicated all data is backed up offsite every 24 hours.

7 Sourcecode

Vergic uses a hosted solution for handling our source code. Vergic uses Visual Studio Online where we have a private, cloud-hosted code repository for TFS and Git.

8 Hosting options

Hosting of the Vergic Engage solution is offered through three (3) different options. Type of hosting option has an impact on the price.

8.1 Vergic Shared Tenant

- Vergic Engage multi-tenant solution
- No setup required
- Traffic goes via https internet connections directly to the Vergic Cloud Computing Platform in Sweden
- Optional – dedicated IDP, Log in server, non-public

8.2 Dedicated Hosting

- A dedicated privately reserved tenant in Vergic Cloud Computing Platform
- Traffic goes via https internet connections, or via an optional fiber channel directly from the Swedish Vergic Cloud Computing Platform to the Customer Infrastructure. VPN Connections can also be offered
- Optional - Private VPN or fiber channel

8.3 Dedicated identity provider

In cases where user authentication need to be further restricted and controlled Vergic offers a dedicated identity provider server. The Vergic identity provider is totally separate from the Vergic Engage application and is hosted in a different tenant than the Vergic Engage application. Any number of dedicated instances of the Vergic identity provider can be installed and limited number of customizations can be offered. A dedicated identity provider tenant also offers the ability to restrict user access, i.e. to certain IP addresses or IP ranges.

All user authentications towards Vergic Engage are done through the SAML2 standard.

8.4 On Premise Alternatives

- Worker nodes are physically located in the enterprises own hosting centers
- A VPN, fiber channel or similar is used to administer worker nodes (updates, monitoring, scaling)
- Worker nodes are centrally administrated from Vergic Cloud Computing Platform
- Traffic is routed through enterprise firewalls
- Required hardware and cost – See separate documentation. Fiber channel / VPN can be offered separately
- Additional custom configuration might occur, based on customers hosting infrastructure complexity and demands