Community
Creation
Commerce

Artwork by Anshe Chung Studios

Volume 2, Number 3
Technology, Economy, and Standards
October 2009

Sponsored in part by:

**The Journal of Virtual Worlds Research
is owned and published by:**

TEXAS DIGITAL LIBRARY

department of
RADIOTELEVISIONFILM
UNIVERSITY OF TEXAS AT AUSTIN

SIRC THE SINGAPORE iNTERNET RESEARCH CENTRE

virtual worlds research consortium

Content Level Gateway for Online Virtual Worlds

By S. Van Broeck, M. Van den Broeck and Zhe Lou
Alcatel-Lucent, Belgium

Abstract

This paper focuses on protecting future virtual worlds from the familiar but often irritating spam, pop-ups, adware and other web-based nuisances.

Over the last decade, the internet had a profound effect on how we work and how we arrange our personal and professional lives. Along with the many advantages, people are also burdened with inefficient parental control, spam, pop-ups, viruses, adware, spyware and more. Virtual worlds available to the general public will obviously be accessible via the Internet, just like the Web pages people are visiting today. Virtual worlds will therefore also be faced and have to deal with similar dangers and negative influences that users and administrators are experiencing today but now manifesting themselves as 3-D models, avatars, textures, animations, or any other type of content commonly used by virtual worlds. We propose a solution to guard our future internet from such counterproductive content.

Keywords: Policy control; virtual worlds; MPEG-V; spam; OSG.

Content Level Gateway for Online Virtual Worlds

By S. Van Broeck, M. Van den Broeck and Zhe Lou
Alcatel-Lucent, Belgium

Along with the many internet advantages, people find themselves overwhelmed with inefficient parental control, spam, pop-ups, viruses, adware, spyware and other web "junk." All these elements are working against productivity and against enjoyment of the internet and need to be counterfeited by specialized software like spam filters, firewalls, adware filters and others. Market revenues in this area reach in the billion of dollars.

Virtual Worlds (VW) are posing themselves as the future of internet. Today, hundreds of VWs already exist, each addressing a target group. As such, there are VWs that are specifically created for educational and training purposes, others focus on travel, social networking or gaming, and still others target communities like corporate environments or young children. It is this mix of VWs and the seamless interoperability between them that represents the three-dimensional (3-D) internet of tomorrow. People in the future will be able to go into virtual showrooms, watch the car of their dreams in their favorite color, enter inside, activate the controls, talk to other visitors, and make a deal with a sales person. They will be able to make a virtual trip, take personalized courses with hands-on exercises on virtual models, and have meetings and perform tasks in virtual settings that are not possible in the real world.

These VWs will have to deal with the same negative influences as the internet of today, only now these distractions will be presented in a different package. Such packages can be 3-D models, avatars, textures, animations, scripts, sound, or any other type of content commonly used by virtual worlds. People will be subjected to aggressive, violent, annoying, sexual, intimidating, or other offending content.

This paper proposes a solution to guard our future internet from the very start from such counterproductive content.

Related Works

Current filtering systems operate on specific protocols including Hyper Text Markup Language (HTML), Simple Mail Transfer Protocol (SMTP), and Internet Message Access Protocol (IMAP). Several standardization initiatives are addressing this topic by defining vocabularies, categories, or semantics describing types of content. As such, the World Wide Web Consortium (W3C) maintains the Platform for Internet Content Selection (PICS) recommendation for HTML content and the Family Online Safety Institute (FOSI) maintains the Internet Content Rating Association (ICRA) standard. Other organizations make use of these standards to filter content based on policies set by a central or supervising authority. As such, companies can implement a company-wide policy for access to HTML content. One such policy control is the eXtensible Access Control Markup Language (XACML) from the Organization for the Advancement of Structured Information Standards (OASIS).

Certain content, however, may not be tagged or may be tagged incorrectly. Therefore, a number of complementary actions can be taken, including black and white lists, content analysis

optionally complemented using data mining or machine learning techniques, statistical data compression models, or user feedback statistics. A lot of effort is spent to identity every single peace of content on the web and to prevent inappropriate content from reaching certain target groups.

To some extent, mentioned techniques can be re-used in VWs. These can, for example, bar certain users from accessing a particular VW much like barred access today to certain internet sites. This solution does not, however, provide a fine-tuned control where an authority is able to set a policy on certain content within the VW instead of on the complete VW. Additionally, the structure of a VW is quite different from the structure of a web site and therefore requires also adapted techniques.

Today, most VW have specific solutions build into their VW platform. As such, Linden Lab's Second Life is using a combination of (1) access rights on object level and (2) categories on land or parcel to refrain certain individuals from accessing certain content. In Instant Messaging Virtual Universe, persons with an adult pass (AP) can see naked persons while persons without such AP will see them with standards clothing. Such solutions are platform specific, do not address interoperability between different platforms, and therefore do not support 3rd party policy control management over content spanning several VWs.

Research has been undertaken to define methods tailored to screening certain content for a specific environment. Such methods can be used to modify or replace certain objects for certain users. As such, a religious person may be presented with a decently dressed person instead of the scarcely dressed one. These are however also specialized solutions often incorporated in the VW platform itself.

System administrators having to deal with all these VWs will either have to deal with the specifics of each of these platforms or have to ban access to the complete platform. In the internet of tomorrow, where numerous VWs will need to interwork, a more fine-grained solution is needed where an authority can create a single policy for content handling applicable to all VW platforms.

Solution

VWs consist of content that is inherently different from the content filtered by current filtering applications. Billboards may have textures that show violent behavior, models may have sexually offending animations attached, and automated avatars or bots may annoy people by offering all kinds of merchandise. To safeguard people from these disturbances, a policy mechanism should be put in place that can be configured by a central authority on behalf of individuals or organizations and that is applicable for all VWs. Similar to current internet policy management, we propose to locate this functionality in access equipment like routers and gateways.

Figure 1: An example of an offending texture within a virtual world.

In order for such policy to work, every type of content should be labeled. Such labels can be part of the content alike the PICS format, but they can also be located separate from the content in semantics alike the ICRA format. As an example of the first one, the extendable COLLAborative Design Activity (COLLADA) format could be extended to hold PICS information on a per element basis. The second approach implements a separate set of semantics, using technologies such as the Resource Description Framework (RDF) language, that link uniquely to the content and that can reveal information about the content. In either approach, models, textures, animations, and any other elements can be labeled with the type identifier for use by the policy system.

Most, if not all, VWs are organized in a hierarchical way. One such library providing a hierarchical structure is the OpenSceneGraph (OSG). Access to the element's hierarchy level as programmed in the OSG could also be used by the policy control. Indeed, once an element is encountered with content type that is rejected by the current policy, all dependent elements can also be refused. The COLLADA format could be extended or the streaming protocol updated to include the current hierarchy level of the element.

Once the type of content can be discovered by a central authority, policies can be set up in a Policy Management Point (PMP) to intercept certain content according to the user's preferences. In case of interception, a specialized application or rule engine can decide on what actions to take next.

In the figure below, a functional overview of the solution is given. In the drawing, five different options for countermeasures are given in case a policy rule is violated and therefore a non-compliant condition encountered.
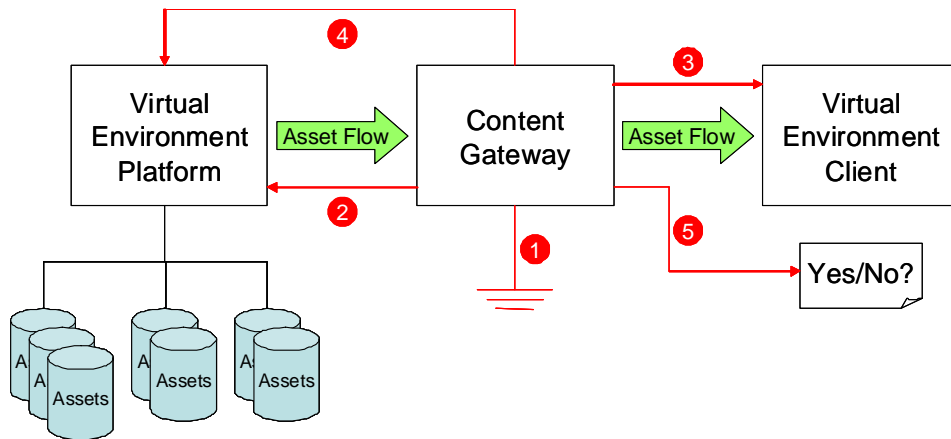
Figure 2: A functional overview of content control.

One of the actions that can be taken by the gateway (arrow 1 above) is to simply abort the application, to discontinue all further content, to remove all intercepted content, or to remove all intercepted content and all their dependents. Simply removing objects from VWs may lead to conditions where the VW scenery becomes incomplete, where the application logic may not be able to function correctly anymore, or where inconsistencies are created between different clients that have a different set of policies for the same VW. The religious person mentioned previously may thus see the scarcely dressed person differently from other people.

In order to overcome such barriers, adequate autonomous removal or replacement of elements requires quite some understanding of the VW platform and application logic. To remedy the simple approach described above, autonomous backward communication with the VW platform presents another viable option. Arrow 2 illustrates this behavior. As such, the VW platform can take corrective actions like sending back adapted content and also distributing these corrections to other clients. When a child enters a virtual room, all offending content may as such be replaced instantly for all the people already present in that room.

Intelligence in the gateway may also autonomously replace content received from the VW platform into other content. For instance, as indicated by arrow 3, non-compliant textures may simply be replaced by a specific and well recognizable default texture much similar to the text a browser shows when a non-compliant page is opened.

As indicated by arrow 4, intelligent VW platforms may also decide to first retrieve the policies in place for a certain client from the Policy Information Point (PIP) so that it can already adapt the VW in accordance with the policy settings. In such case, the intermediate Policy Execution Point (PEP) will continue to operate on the incoming content but will most probably never need to intercept. This third approach also allows every individual VW application's logic to implement corrective measures that best match their logic. As such, one VW platform may, for example, decide to disallow the user from visiting certain places while another VW platform may choose to simply substitute the violating content with an acceptable alternative.

As for current existing solutions, the described policy control can be extended with black and white lists for VW platforms as well as for individual elements within the VW, user feedback on VW platforms and elements (see arrow 5), algorithm-based content inspection, statistical information, or any other existing means to help identity the type of content. For example, textures may be screened to find nudity, models and animations can be analyzed in search of obscenity, and scripts can be evaluated to discover misbehavior.

It may prove beneficial to introduce a general replacement strategy for elements that get rejected by a policy and where the VW platform can or does not take corrective actions. In case of animations, the violating animation could be replaced by an animation that reflects refusal. In case of textures, a generalized texture could be defined indicating the violation to the user.

In order for the policy to execute the screening of the content, it must have guarded access to the content. Therefore, it may be necessary for VWs to agree on security measures with the policy execution point. Both parties may use a Public Key Infrastructure (PKI) to secure their communication or decide to simply make use of secure protocol layers like Secure Sockets Layer (SSL).

Conclusion

To safeguard the future of the internet, where many different VWs will co-exist and will have to interoperate with each other, a policy based control is needed to safeguard the users from unwanted or malicious content. This paper proposes a solution where all VW content can be labeled and based on this label, screened by a policy authority. In case of violations: (1) several simple autonomous corrective measures that can be taken, (2) a communication means with the violating VW to allow the VW to take corrective measures, and (3) a communication means for the VW to anticipate violations by retrieving the policy before streaming. On top of this policy-based mechanism, non policy-based existing measures and means to block unwanted content remain applicable. This paper has been written in scope of the Information Technology for European Advancement 2 (ITEA2) Metaverse1 project that is in charge of MPEG-V (Moving Pictures Experts Group for Virtual Worlds) standardization.