Multi-domain Lightpath Authorization using Tokens.

Leon Gommans, Li Xu ,Yuri Demchenko, Alfred Wan, Mihai Cristea, Robert Meijer, Cees de Laat
University of Amsterdam

Abstract.

This paper highlights the concepts and results of our research leading to demonstrations during the period 2005-2007 to develop a flexible and simple access control model and corresponding support tools to provision multi-domain optical network resources on demand. The paper introduces the general network resources provisioning model that extends the Generic AAA Authorisation sequences for multi-domain scenarios and explains how token based access control and policy enforcement can be used during the provisioned resource access. To build a solid conceptual foundation for the proposed token based access control, the paper revisits existing token definition and proposes a new definition in the context of our research (this paper). The paper subsequently explains the use of tokens during different stages of the lightpath provisioning process. The paper identifies and describes two major scenarios in multidomain lightpath provisioning: The chain and tree approaches that correspond to federated provider based multi-domain resource management and centralized management typical for current Grid based resource management. The proposed token concept allows simple combination of the access control enforcement at different networking layers: the packet layer, the path layer, and the service layer. The paper provides a brief description of a few experiments and demonstrators that proves the proposed concepts and solutions that illustrates its acceptance by a wider networking community.

1.0 Introduction

Modern high performance distributed applications, dealing with high volumes of data, increasingly require dedicated high-speed optical network connections that are provisioned in an on-demand fashion. This type of resource is commonly referred to as a lightpath[1]. Projects, such as OptIPuter[2], envisage a LambdaGrid, where lightpaths are tightly coupled with computational resources. A LambdaGrid coordinates dynamic provisioning of end-to-end circuits using Grid concepts. On the other hand, large Grid projects such as the LHC Computing Grid[3] use their own dedicated network infrastructure, designed to handle the required data volumes without being tightly coupled to computational resources. In our paper we will not target such applications but consider data intensive applications that are expected to benefit from the ability of a network to dynamically allocate and reserve lightpaths that are shared at different times with other applications. Several examples of these applications within areas such as data mining and visualisation can be found within the realm of the OptIPuter project. We will also consider network situations were multiple network providers must work together in order to create end-to-end lightpaths. We will assume that providers will allow applications or its middleware to make lightpath reservations. As lightpaths typically do not use network layer data forwarding techniques and rely on layer-2 or below technologies, access control to a lightpath becomes more difficult, in particular if a lightpath needs to be specifically bound to an application. During the course of this paper we will see that network domains and applications can work together in different ways to make sure applications, which reserve a lightpath, actually get unique access to their reserved lightpath.

Hybrid networking concepts within networks such as SURFnet6[4], Internet2s Dynamic Circuit Network[5] (DCN), CA*Net UCLP[6], G-lambda[7] and GEANT2 Autobahn[8] allow applications to reserve and use a lightpath on demand. Within these networks it is however unclear how particular applications can be given exclusive access to a reserved lightpath, whilst preventing other applications from using the same lightpath during its use. In this paper, we show a token based access control mechanism that can be used for this purpose. Recent research and development projects such as Phosphorus[9] and Internet2 DCN aim at making network resources Grid middleware enabled. The token approach is being incorporated and tested in these projects.

A token provides a flexible mechanism that allows the right to access a lightpath to be

associated with a request from an application. After a user (or application) requests access to a network resource, the network is capable to recognize a token that enforces access across multiple domains. We will show how tokens can prevent other users or applications from gaining access to the same resource at the same time. The focus of this paper will be on the access enforcement ability of the network, its granularity, and ways how the network can create the associated context needed to enforce a token. We will only mention some of the policy-based decision types that domains typically make before they decide to grant access.

In the paper, we will first elaborate on the concepts around tokens. We will then briefly describe how these concepts were applied in various provisioning and access control enforcement models. We will end by briefly describing demonstrations during subsequent iGrid 2005 and Supercomputing (SC) events in 2005, 2006 and 2007.

2.0 The token as a concept in networking.

In this chapter we will elaborate on the concepts around tokens in the context of networking. Questions like "Why use tokens?", "What is a token?", and "How are tokens created and handled?" will be discussed.

2.1 Why use tokens?

Current optical network control and management plane implementations do not employ mechanisms that consider and enforce data-flows from individual application sessions. These implementations enable users to reserve and allocate a lightpath. After allocation, the application signals the network using protocols such as RSVP-TE[10] or XML/SOAP that it likes to use the lightpath. The allocation typically specifies a lightpath between two endpoint addresses, for example physical port numbers or IP addresses. The network typically assumes that the application component is directly connected to the specified ports. Most mechanisms will first authenticate and subsequently authorize the application user before allowing the user to make a reservation for time and bandwidth between the endpoints. Once completed, the network does however not enforce the relationship between the user dataflow and the lightpath. The network assumes that the application will use the same ports as requested. It also assumes that no other applications will share the connection at the same time. These assumptions make authorized public usage of hybrid networks, offering lightpath services, more complicated. In addition, authorized usage becomes more complex when the reservation process involves multiple domains. In such case, the downstream domain must trust the upstream domain that it forwards the intended flows. The pictured problem is not unlike making reservations on a multi-legged flight and selecting seats, without the presence of airport authorities and/or airline employees to enforce access to the intended plane and its seat. Without such enforcement, anybody could board the plane and occupy the reserved seat without the rightful person being able to prove his/her right to be seated on this flight. Airlines use boarding passes. In networks we propose to use tokens for the same purpose.

2.2 What is a token?

The word "token" is an overloaded term. The term is likely to create confusion if we do not define it in the context of our research and paper. While the generic meaning of the word "token" is "a visible or tangible representation of something abstract", "a characteristic or distinctive sign or mark", the "security token" as it is defined in the Web Services Trust[11] context actually means a security protected credential. Within the context of this paper, we therefore use the following general working definition for a token:

"A shared abstract permission that is presented as part of an access request in each domain"

The permission is a small piece of information that unambiguously references information providing the context of a specific lightpath session. Tokens are used as part of a security scheme, where its possession proofs a right when challenged during resource access control phase. Tokens are different from certificates and tickets in the sense that a certificate carries multiple attributes in a specified format and each attribute has a defined and explicit meaning. A ticket also carries attributes but its scope and validity is limited and its format is application dependant. Tokens, certificates and tickets have in common that they are integrity protected and its authenticity is ensured by the issuer or signer. In comparison to a certificate or ticket, the meaning of a token is strictly abstract.

A token is obtained, carried and presented by a holder. The recipient must understand its abstract meaning. This understanding may be contained in the logic of recipients program and may be augmented by the authority before a holder presents a token. The same token may express different meanings when the holder presents it to multiple recipients. Authorities must therefore make all possible recipients aware of the relevant meaning of the token. This may be perceived as a disadvantage, however tickets or certificates recognition by multiple recipients require that its attributes must share an agreed meaning. The abstract nature of token allows flexible usage in multi-domain lightpath provisioning scenarios. A token references a shared, context dependent meaning.

2.3 Authenticity of a token.

A token must carry a proof of its authenticity. This can be achieved by using a secure message authentication algorithm (e.g., HMAC-SHA1) to calculate (part of) the content of the token that must be recognized by the recipient. The key, used in the algorithm, must either be shared between authority and recipient or the recipient must have an exact copy of the token. In this way a trust relation will be established between authority and recipient. If the digest used to generate and verify the token includes (part of) the service related context, the user will not be able to modify this context without invalidating the token. We will see later that the token can be used at different layers. At the IP layer, the token digest can for example include the IP addresses, TOS value, etc. Modifying the destination IP address of the packet, will invalidate the token. We will see that higher layers typically use a unique session ID as digest.

2.4 Tokens as part of an authorization sequence.

The presented solution is based on the further development of the AAA Authorization Framework RFC2904[12]. The push model, described in this framework, has been used in scenarios that implements network resource provisioning involving multiple domains. The provisioning process can be split into three stages[13]: (1) reservation / authorization, (2) deployment or activation, and (3) access or use/consumption. The reservation stage, which involves the user, may require (sometimes complex and time consuming) interactions to find, select, schedule and authorize the appropriate resources. In section 5 we explain that our implementation allows authorisation languages such as XACML[14] and SAML[15] to be used during these interactions. We will assume that resources can be committed after relevant authorization decisions have been made. Subsequently we assume that the reserved resources can be associated with a common access control token at the end of stage 1. During stage 3, the token will be presented as a part of the network access request in each domain. At this stage, a token will be evaluated against the reservation context (meaning) stored during phase 2 inside a domain that is referred to by the token. Fig 1 illustrates the extension to the RFC2904 push sequence for a token. The addition to this sequence is the part where the token meaning is provisioned by the authority. Note that, as its meaning is explicit, this part may not be necessary in case the authority replies a certificate or ticket.

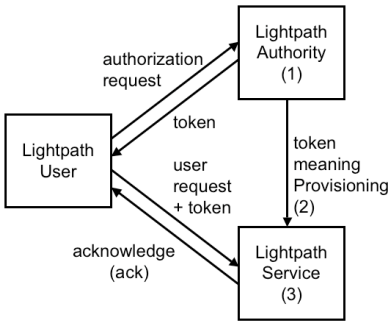

Fig. 1. The basic token sequence as an extension of the basic RFC2904 push sequence showing the position of the three provisioning process stages.

Also note that fig. 1 only shows the interactions needed to communicate authorization, not the actual use of the lightpath by the user.

The above sequence is aimed at allowing a Lightpath Authority to be flexible in assigning a specific context to a commonly agreed token. The Lightpath Authority is involved in the reservation/authorization decisions made during stage 1. The deployment stage (2) performs token meaning provisioning where the reserved resources are typically bound to some reservation ID carried by the token. We will refer to this ID as the Global Reservation Identifier

(GRI) that will be described in more detail later. The Lightpath Service performs stage 3. Stage 3 is like checking the passenger boarding the plane. The possession of a token enables the passenger (i.e. a user accessing a lightpath segment) to be checked whilst boarding the plane. When checking in on the next leg, the same token containing the reservation number (playing the role of GRI) can be used to refer to a different "seat number" (the context describing the next lightpath segment). This brings us to the subject of multi-domain scenarios

3.0 Tokens in multi-domain scenarios

Here we consider the role of a token during the handling of a request by authorities in a multi-domain scenario leading to stage 2. We will then look at how tokens can be enforced inside the service entities at stage 3. To allow multi-domain lightpath provisioning, the domains must interact in a coordinated manner. Here we distinguish two typical approaches: The chain and tree approach. The chain approach is typical for multi-domain network provisioning scenarios used amongst Network Service Providers. An example.of this approach can be observed within Internet2s DCN network where InterDomain Controllers (IDCs) operate as domain Lightpath Authority. We will elaborate on this scenario in section 5. In chapter 3.3 we will discuss the tree approach, typical for Grid scenarios. We will first discuss the chain approach.

3.1. Context provisioning and token creation using the chain approach.

When a user during stage 1 requests an authorization from a Lightpath Authority to use a particular lightpath in a typically multi-domain optical network, each domain's authority will apply some policy when evaluating a request. Policies may imply rules and/or conditions regarding the identity of the requestor, its authorizations, the existence, route and (optimal) availability of the requested path, priority of the request, etc. Each domain may have its own policy what will imply a specific domain related context to a decision that the token will represent.

Fig 2 illustrates interactions between major entities participating in a multi-domain lightpath provisioning chain approach scenario. The process is initiated by a user request sent to the domain A's Lightpath Authority. At this stage, a GRI is created by domain A. The GRI, must be a globally unique identifier. It can be either be implemented as a, large, randomly generated number, that can be considered as sufficiently unique, or as a domain-unique number concatenated with a unique domain identifier.

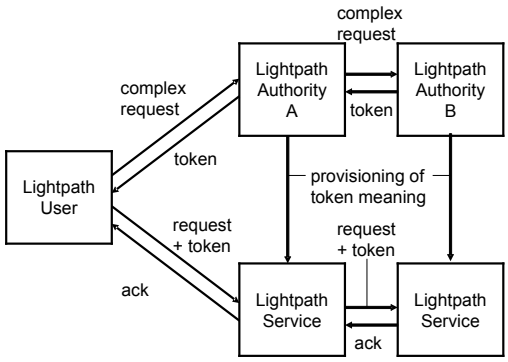

Fig. 2. Provisioning a multi-domain chain of domains.

The GRI serves to identify a lightpath session across multiple domains. The GRI may also be used inside a domain to administer local resource details. The outcome of the policy decision process is either positive or negative. The negative result is logged and replied to the requester. A positive result will cause a request to be administered and sent to the next domain (B) along the path. This request will include the GRI. A subsequent decision taking process may again yield a positive or negative result. The GRI is used to administer the result and its details in domain B. A negative result is returned to the upstream domain (A). A positive result at this stage means that all previous domains can serve the request. Being the last domain in the chain, it also means that the entire request can be honored. As a way to express this fact, a process in the last domain will create a token by applying a secure message authentication algorithm (HMAC), to create a digital signature from the GRI using either a shared secret or trusted key of the last domain. To simplify secure context management, the token might just consist of the GRI and its signature. If not mentioned differently, we assume such a token in the remainder of the article. The signature might however be produced, by including part (or all) of the reservation context into its generation process for reasons discussed in 2.3. This step concludes stage 1 and will be followed by stage 2, where the reserved resource deployment / activation takes place. Stage 2 essentially means

provisioning each domain Lightpath Service with the token or token-key and its associated meaning. This information allows token recognition and verification at the resource access stage (stage 3). At stage 2, the Lightpath Authority of domain B provisions the Lightpath Service of its domain. Lightpath Service B can use the GRI as an index to store the characteristics of the lightpath (bandwidth, time, ingress / egress points, etc.) Domain B must also, at end of stage 1, return the token or the key used to generate the token in the reply to domain A. Domain A will administer the reply, using the GRI as index. This will then also enable the service part of domain A to be provisioned.

The user will now receive from domain A the reply that includes the token (containing the GRI). At the agreed time, the user will signal the lightpath and include the token in the request. By comparing the token with the provisioned token (either provisioned directly or re-generating the token using the provisioned token-key), the Lightpath Service can quickly verify the validity of the token and provision the requested circuit. The GRI part of the token can be used to lookup the corresponding reservation context and token/ token key that can be used for token validation. The request is then forwarded to the next domain where the same token is used as a means to perform access control to a set of different resources indexed by the GRI of the token. Note that communication during stage 1, 2 and 3 may be secured using a shared secret model or use a PKI based inter-domain trust infrastructure. This kind of security is considered independent of the security used to make the GRI authentic, i.e. creating the token.

3.2 The token context.

As discussed, each domain in fig 2 may associate a different meaning or context to a token: The token may refer in domain A to bandwidth for a specified amount of session time between a specific pair of ingress- and egress ports. The information about ingress and egress ports will be different for domain B. Moreover, domain A may use a different time slot granularity than domain B. If A uses 1 minute timeslots and B uses 5 minute timeslots, then allocating a 12 minute lightpath means 12 minutes in domain A, but may be translated to 15 minutes inside domain B. Allowing authorities to each provision a different service context to a token is an essential characteristic.

3.3. Multi-domain context provisioning and token creation using the tree approach.

In Grid environments, the network resource may be provisioned in the same way as any other Grid resource. Grid applications typically use a centralized scheduler as common authority for this purpose.

Fig. 3 shows the tree approach. In Grid environments resource reservation and scheduling is a part of the middleware functionality. In collaboration with Santa Clara University, University of Amsterdam investigates the use of an elastic scheduler[16] to reserve network resources. Part of the Phosphorus project researches the functions of the ISS/VIOLA[17] meta-scheduler for finding optimal choices when co-allocating network- and computing resources involving multiple domains. Additionally, combined grid-network resources reservation may allow creating optimal mapping between grid jobs and required distributed computational resources with network performance limitations. This is subject of ongoing research in the G-Lambda[18] and Phosphorus projects.

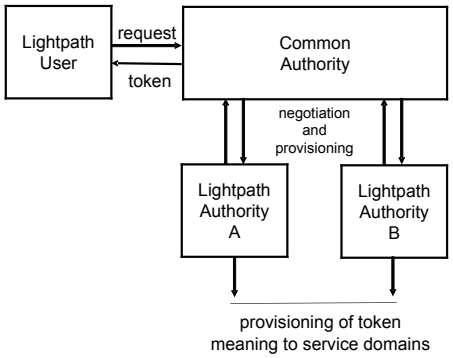

Fig. 3. The tree approach.

Within the tree approach, a common authority will negotiate with individual lightpath authorities along the path. If the common authority can resolve the request, it will provide a token to the user to indicate all involved domains are committed to provide the requested resource. Alternatively, each domain can create a token, where the common authority just passes it on to the user. In this case, the user needs to insert a number of tokens into the signal to use the lightpath, one for each domain. The feasibility of using the same provisioning and

policy enforcement model for both approaches is part of our current research. We expect that tokens and the concept of a GRI can glue together both chain and tree style authorization.

4.0. Access control granularity and enforcement layers

After the context is provisioned and stored inside the Lightpath Service, the service will wait for a service request to arrive for subsequent enforcement. When received, the GRI part of the token points to the context of the lightpath reservation stored by the Lightpath Authority. The service request with the token can be sent in a number of different ways:

1. **At IP packet layer.** Each IP packet is considered as an individual access request. At this level, the token is included inside each IP packet, e.g. inside the IP Options field of IPv4 packet. This enables per packet enforcement. As per packet access enforcement is common in firewalls, we call this approach the firewall- or packet layer approach. At this level the token is typically a secure hash result of the context (e.g. content of IP packet header) and may even not contain a GRI.
2. **Control plane or Network layer path signaling.** Each path-signaling message, such as an RSVP-TE PATH message, contains a token. As RSVP-TE messages are sent at certain time intervals to keep the data-path alive, this kind of signaling will enable enforcement by keeping the path alive. An invalid token could cause a teardown of the path or could stop the forwarding of RSVP-TE messages by a Label Switch Router (LSR). Tokens could be placed inside a Policy_Data object as defined by RFC2750[19]. We call this approach the path signaling approach.
3. **Service layer signaling.** Service layer signaling typically employs an XML based protocol such as SOAP to implement a Web Service. A token can be part of the object exchange. The service application logic will determine if a single token exchange is sufficient to authorize the resource access or that a token must be sent periodically to keep the circuit alive. We call this approach the service layer approach.

Note that each of these different approaches implies different levels of enforcement granularity. At IP packet layer, we have the finest granularity where each packet is subject to access control, whereas the approach at service layer could only be enforced once, i.e. when a lightpath is signaled when connecting.
Examples of these approaches were shown during subsequent Supercomputing events of 2005, 2006 and 2007 and during iGrid2005.

5.0 Implementation of the Token based access control method.

The token based access control mechanisms have been implemented as components of a general authorisation infrastructure for network resource provisioning. It is used to simplify access control to reserved distributed resources in a multi-domain environment. The infrastructure, called the Generic AAA toolkit (GAAA-TK), is being developed by the University of Amsterdam (UvA). The GAAA-TK both implements a number of security mechanisms to support the multi-domain policy based authorisation process as well as the token-based access control. As such, the Token Validation Service (TVS) has been developed as a special component to support token handling at all stages of the general network resource provisioning. It supports interdomain token based signaling during the reservation stage. It performs path and reservation context distribution at the provisioning stage. It also provides the token validation service at the access phase. The GAAA-TK is provided as a pluggable Java library and as a standalone domain central authorisation service (DCAS). The special GAAA-TK profile and TVS implementation includes support for all layers mentioned in chapter 4.0. The GAAA-TK also implements the SAML-XACML[20] authorisation request-response protocol that allows for authorization request evaluation with the local or remote XACML based Policy Decision Point (PDP).

Although the TVS component has been implemented as a part of the general GAAA-TK library, it can also be used separately. All basic TVS functions are accessible and requested via a Java API. As such it can be used with other authorisation services implementations and frameworks such as Globus Toolkit Authorisation Framework[21] and PERMIS[22] to support necessary functionality for token distribution and processing in their target application areas. Further TVS development will

extend Web Services interface to allow all TVS functions be accessible via Web services. The current TVS implementation supports both shared secret and PKI based token key distribution.

6.0 Demonstrations of the token principle.

In this section we will present some of the work that has been done within the context of projects that collaborate and share information directly or indirectly with the OptIPuter project. Various aspects of the tree and chain token approach where demonstrated at different occasions.

6.1 The packet level approach.

The packet level approach was demonstrated using an Intel IXDP 2850 NPU development platform programmed as Token Based Switch (TBS) at SC2005[23]. Here a token, inserted into the IP Options field, enabled IP packets to take a specific pre-provisioned lightpath. This offers IP layer support at stage 3. For implementation details we refer to [24]. OGF document GFD.083 Firewall Issues Overview[25] argues that this kind of switch could form a potential solution for a firewall, protecting hybrid network resources if public access needs to be supported as mentioned in section 2.1. In later releases of the TBS we programmed it to forward a token received, at IP layer, to the RSVP-TE layer and also XML/SOAP layer, as such acting as a token gateway.

6.2 The path signaling approach.

This approach was the subject of our demonstrations during SC2006. Fig. 4 illustrates its components.

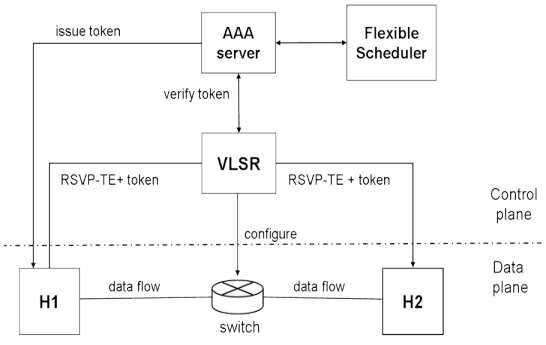

Fig 4. Token-based GMPLS at the path layer.

Here we demonstrated how a GMPLS based network is able to support tokens by including it in a specific field of a RSVP-TE PATH signaling message. To show this ability we modified the Virtual Label Switch Router (VLSR) and Client System Agent (CSA) code of the open source GMPLS project - DRAGON[26] to recognize tokens. In this demo, the mentioned elastic scheduler (chapter 3.3) acted as an advance reservation resource manager to take decisions for stage 1. The token was stored inside the AAA server for stage 2. At stage 3, the tokens were inserted into a Policy_Data object (RFC2750) of RSVP-TE PATH messages that are exchanged between hosts and the VLSR to signal the data-path setup. The VLSR parses the request message and verifies the token by querying the Generic AAA server. If the token is signaled valid, the VLSR forwards the message to the next hop and configures the switch in the data plane.

6.3 The service layer signaling approach.

This approach was subject of a single domain demonstration during iGrid 2005 and Supercomputing 2005 and a multi-domain case during SuperComputing 2007.

6.3.1. Single domain case: The VM migration experiment.

In our Supercomputing and iGrid experiments in 2005, we used a Generic AAA server from the GAAA-TK as Policy Decision Point (PDP[27]). The Generic AAA server architecture is described in more detail by RFC2903[28]. Fig. 5 shows the basic setup of the experiment performed during iGrid 2005 and SC2005.

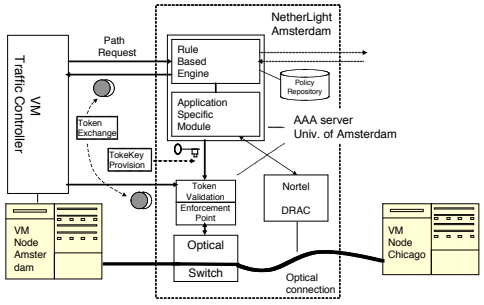

Fig. 5. The VM migration experiment.

The experiments[29,30] were conducted in collaboration with Nortel. Here we showed the migration of a XEN Virtual Machine (VM)

across a lightpath. A Micro Electro-Mechanical Systems (MEMS) based Optical Switch enforced access to the lightpath by switching an authorized CPU of a cluster to the designated lightpath. Nortel's Dynamic Resource Allocation Controller (DRAC) was in control of provisioning and resource management of a lightpath between Amsterdam and Chicago. In the demo-scenario, a VM Traffic controller wanting to migrate a Virtual Machine via a given lightpath initiated stage 1. After contacting the DRAC to check if the request can be honored, the Generic AAA server (consisting of a Rule Based Engine and Application Specific Modules – see RFC2903) generated a token. During stage 2, the Generic AAA server provisioned the Token Validation (Policy) Enforcement Point with the token-key. At the appropriate time, i.e. at stage 3 when the actual migration of the VM is about to happen, the VM Traffic controller will insert the token into the Token Enforcement Point. If the token is accepted this function will control the Optical Switch such that it will connect the right VM node to the right optical path. The DRAC was assumed to provision the circuit at the agreed time. The mechanism will prevent different VMs from migrating at the same time using the same resource. This example shows that applications can be more accurately associated with a lightpath as stated in section 2.0.

6.3.2 A multi-domain case: Implementation of the Token Validation Service

At Supercomputing 2007 we proposed and implemented the token concept into the IDC control plane of Internet2 DCN. The earlier mentioned Token Validation Service (TVS) was integrated with the IDC, as shown in fig. 6.

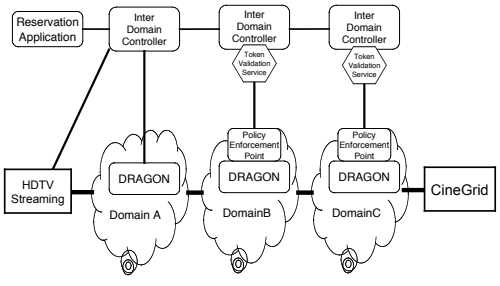

Fig. 6. The Token Validation Service experiment

The TVS enables an IDC to generate and communicate tokens much in the same way illustrated in fig. 2. In the above example a reservation application obtains a token from the chain of IDCs in the same way as described in 3.1. In the scenario we developed for SC2007, a token was subsequently placed on a USB memory stick and carried to a MacMini with a Full HD-TV display to show a movie streamed from a CineGrid[31] server in Amsterdam via a 1 Gbps DCN link. The difference with the previous examples is that Internet2 implemented an IDC version where tokens were handed back to the Lightpath Authority at stage 3. This was considered the easiest solution for a first implementation. The path signaling way, using a LightPath Service implemented with the GMPLS implementation from the DRAGON[32] project together with a Policy Enforcement Point developed as part of the TVS at UvA, was implemented at a UvA testbed. Also, not all domains may want to support token enforcement. Fig 6 shows domain A without such ability. For such cases, the inter IDC protocol supported transparent pass-through of tokens. After a reservation is made and the enforcement points are provisioned (stage 1 and 2 complete) the IDC is signaled to open the reserved path for stage 3 using only SOAP/XML messages.

7.0 Future Work

Combining the use of tokens with the tree and chain signaling approach, that use the same interface for a domain will be a key item for further research. This should enable domains to authorize the use of network resources to create a lightpath in many different scenarios. Also some form of GRI and token format will need to be agreed upon, such that domains can identify lightpath sessions and base their internal administration and enforcement on it. Work on this within the GLIF and research projects such as Phosphorus and GigaPort is currently ongoing. In a simple scenario, the TVS can be programmed with a shared key using a Web Services interface to facilitate communication between the Lightpath Authority and Lightpath Service. A more flexible and automatic TVS security model may use an Identity Based Cryptography[33] (IBC) approach that relies on a domains IBC key generation service.

8.0 Conclusions.

The paper presented the results of our ongoing research and development to build a consistent authorization architecture and flexible access control infrastructure for multidomain hybrid network provisioning. The proposed and discussed concepts and solutions use a common abstract token concept.

We have shown that a token can act as a *shared abstract permission that is presented as part of an access request in each domain* where its permission is represented as an index pointing at a pre-allocated network resource. In multi-domain scenarios the same token may point to different definitions of a lightpath segment inside different domains. As such a token can be considered as glue to collect authorizations to use network segments inside different domains, forming an end-to-end lightpath.

We showed how tokens are used at all three stages of the RFC2904 based resource provisioning sequence: The access token is created as a result of the successful phase 1 during which the multidomain path is reserved. During the following phase (2) the reservation and token context information (including the token key) is provisioned to all participating domains. In the following lightpath access phase (3) the token is used to enforce access to the network resource. The abstract nature and small size of tokens allow their use for access control enforcement at three different networking layers: the IP layer, the path layer and the service layer and showed examples of their usage. We also showed that two different models are common during the collection of authorizations to create a token: the tree and chain model.

In our experiments and demonstrators we proved that the token mechanism is a flexible and powerful way to allow different domains to share and enforce lightpath authorizations. We exploited simplicity and flexibility of the token as it can be contained by different protocols and is able to be passed on between protocols. The GMPLS control plane can forward the token inside an XML based messages such as a SAML assertion. Also the fact that the usage of the token is completely independent from the way domains negotiate in either the tree or chain fashion is a powerful concept that facilitates interoperability. A Lightpath Service that enforces tokens does not care how it receives the provision information at stage 2 as described in section 2.4.
Further investigation of these characteristics is a logical continuation path for our research into how domains can interact to offer authorized lightpath services.

We proposed and jointly used the GRI concept as a common session identifier in our collaborative effort with the Internet2/DCN project. The GRI was used as a resource identifier that is created at the beginning/start of the provisioning session/process to simplify the provisioning process tracking. We looked at a signed GRI as a possible form of a token. We found that this format enables each domain to keep administrative details of its lightpath segment hidden from other domains, whilst referring to the same end-to-end path. The token subsequently allows domains to enforce access to its resources without the need for an unpredictable overhead to contact the authority. As such, tokens offer a fast and flexible way to allow different domains to share and enforce lightpath authorizations. In our SC2007 demo we consolidated the token concepts into a Token Validation Service (TVS). The TVS supports token handling at all stages of the general network resource provisioning sequence

9.0 Acknowledgement.

The authors like to acknowledge the contributions of the following people involved in the experiments and discussions around the concepts: John Vollbrecht, Tom Lehman, Chris Tracy, Andrew Lake, Jerry Sobieski, Jarda Flidr, Brian Cashman, Franco Travostino, Inder Monga, Phil Wang, Fei Yeh, Eric Bernier, Paul Daspit, Satish Raghunath, Bram Peters, Erik Jan Bos, Pieter de Boer, Ronald van der Pol, Marten Hoekstra, Jeroen van der Ham, Bert Andree, Paola Grosso, Ralph Koning, Arie Taal, Rudolf Strijkers, Freek Dijkstra, J.P. Velders, Jeroen Roodhart, Gigi-Karmous Edwards, Admela Jukan, Joe Mambretti, Tom DeFanti, Bill Allcock, Sumit Naiksatam. The results obtained were part of the work funded by the Dutch GigaPort RoN project, the EU IST FP6 project NextGrid (contract number 511563) and EU IST FP6 project Phosphorus (contract number 034155).

10.0 References

[1] Customer-managed end-to-end lightpath provisioning, Jing Wu, Michel Savoie, Scott Campbell, Hanxi Zhang, Gregor V. Bochmann, Bill St. Arnaud, International Journal of Network Management, Volume 15, issue 5 (September 2005), pg 349-362

[2] OptIPuter project [Online] http://www.optiputer.net

[3] The LHC Computing Grid [Online] http://lcg.web.cern.ch/LCG

[4] SURFnet. [Online]. Available at: http://www.surfnet.nl

[5] The Internet2 Dynamic Circuit Network. [Online]. Available at: http://www.internet2.edu/dcresearch/index.html

[6] UCLP [Online] http://www.canarie.ca/canet4/uclp/

[7] G-lambda project [Online] http://www.glambda.net

[8] GEANT2 AutoBAHN – general info available via GEANT2 homepage[Online] http://www.geant2.net

[9] The Phosphorus Project. [Online]. Available at: http://www.ist-phosphorus.eu

[10] RFC 5151, Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions, A Farrell, A. Ayyangar, JP. Vasseur, IETF februari 2008.

[11] Web Services Trust Language (WS-Trust). [Online]. Available at: ftp://www6.software.ibm.com/software/developer/library/ws-trust.pdf

[12] RFC 2904 AAA Authorization Framwork, J. Vollbrecht, P. Calhoun, S. Farrrell, L. Gommans, G. Gross, B. de Bruijn, C. de Laat, M. Holdrege, D. Spence, IETF August 2000.

[13] Demchenko Y, F. Wan, M. Cristea, C. de Laat, "Authorisation Infrastructure for On-Demand Network Resource Provisioning", The 9th IEEE/ACM International Conference on Grid Computing (Grid 2008), Tsukuba, Japan, Sept. 29 - Oct. 1, 2008. Accepted paper.

[14] XACML: eXtensible Access Control Markup Language, OASIS standard [Online] http://www.oasis-open.org

[15] SAML: Security Assertion Markup Language, OASIS standard [Online] http://www.oasis-open.org

[16] Elastic reservations for efficient bandwidth utilization in LambdaGrids, S. Naiksatam, S. Figueira. Future Generation Computer Systems, Volume 23 , Issue 1 (January 2007), Pages: 1 - 22

[17] Integration of Grid Cost Model into ISS/VIOLA Meta-Scheduler environment, R. Gruber, V. Keller, M. Thiemard, O. Wäldrich, Ph. Wieder, W. Ziegler, and P. Manneback, Proceedings of 2nd UNICORE Summit 2006 in conjunction with EuroPar 2006, Dresden, Germany, LNCS 4375, pages 215-224

[18] G-Lambda: Coordination of a Grid scheduler and lambda path service over GMPLS, Atsuko Takefusa, et. al., iGrid2005 special issue, Future Generation Computer Systems, volume 22 issue 8, pp 868-875 (2006)

[19] RFC2750 RSVP Extensions for Policy Control, S. Hertzog, IETF January 2000.

[20] Using SAML and XACML for Complex Resource Provisioning in Grid based Applications, Demchenko Y., L. Gommans, C. de Laat. In Proceedings IEEE Workshop on Policies for Distributed Systems and Networks (POLICY 2007), Bologna, Italy, 13-15 June 2007. ISBN-13: 978-0-7695-2767-3, ISBN-10: 0-7695-2767-1. pp. 183-187.

[21] Globus Toolkit Authorisation Framework. [Online] Available: http://www.globus.org/toolkit/docs/development/4.1.0/security/authzframe/

[22] PERMIS Project. [Online] Available: http://sec.cs.kent.ac.uk/permis/

[23] Token Based path authorization at Interconnection Points between Hybrid Networks and a Lambda Grid, Leon Gommans, Cees de Laat, Robert Meijer, IEEE GRIDNETS2005 proceedings, ISBN 0-7803-9277-9.

[24] The Token Based Switch: per-packet access authorisation to optical shortcuts, Mihai-Lucian Cristea, Leon Gommans, Li Xu, Herbert Bos, Proceedings of IFIP Networking'07, May 2007

[25] GFD-I.083 Firewall Issues Overview, Ralph Niederberger, William Allcock, Leon Gommans, Egon Grunter, Thijs Metch, Inder Monga, Gian Luca Volpato, Christian Grimm. Open Grid Forum, August 2006.

[26] DRAGON: A Framework for Service Provisioning in Heterogenous Grid Networks, Tom Lehman, Jerry Sobieski, Bijan Jabbari, IEEE Communications Magazine, Volume 44, Issue 3, March 2006

[27] RFC 2748 The COPS (Common Open Policy Server) protocol, D. Durham, J. Boyle, R. Cohen, S. Herzog, R. Rajan, A. Sastry, IETF January 2000.

[28] RFC 2903 Generic AAA Architecture, C. de Laat, G. Gross, L. Gommans, J. Vollbrecht, D. Spence, IETF August 2000

[29] Token Based Networking: Experiment NL101, L. Gommans, B. van Oudenaarde, A. Wan, C.T.A.M. de Laat, R. Meijer, F. Travostino and I. Monga, iGrid2005 special issue, Future Generation Computer Systems, volume 22 issue 8, pp. 1025-1031 (2006)

[30] Seamless Live Migration of Virtual Machines over the MAN/WAN, F. Travostino, P. Daspit, L. Gommans, C. Jog, C.T.A.M. de Laat, J. Mambretti, I. Monga, B. van Oudenaarde, S. Raghunath and P.Y. Wang , iGrid2005 special issue, Future Generation Computer Systems, volume 22 issue 8, pp. 901-907 (2006)

[31] CineGrid Project. [Online]. Available at: http://www.cinegrid.org/

[32] DRAGON project. [Online]. Available at: http://dragon.maxgigapop.net/twiki/bin/view/DRAGON/WebHome

[33] Identity-Based Cryptosystems and Signature Schemes. Adi Shamir, Advances in Cryptology: Proceedings of CRYPTO 84, Lecture Notes in Computer Science, 7:47--53, 1984.