Implementing a VPN service with policy rules

Hanine Abdelkrim & Noël Verhoeven
ALCATEL, Route de Nozay, 91460 Marcoussis, France

Abstract: Today's telecommunication service providers strive to drastically reduce the service implementation time in order to minimise the cost. Simultaneously, they want to provide better quality of service. In this context, a solution identified by the Telecommunication Management Forum (TMF) is to propose a completely automated top-down network management process using the customer needs as input.

This paper illustrates how this management solution can be achieved for the implementation of a provider provisioned VPN service (PPVPN): starting from the expression of the customer needs and resulting with the network configuration. This implementation is defined in several steps. First the formalisation of the customer needs is packed within one or several Service Level Specification(s) (SLS). Then these formal needs are translated using specific policy information models. Finally, these models are implemented in order to address the network layer.

Keywords: Provider Provisioned VPN, SLA, SLS, policy-based management, policy information model, policy information base, service network mapping

1. INTRODUCTION

A crucial objective for a service provider (SP) is to have rapid response times to the service requests of its customers in order to be competitive: more services sold, less provisioning effort. This requires automation [1] starting from the business layer down to the network layer.

At the business layer, a SP manages (written) contracts named Service Level Agreement (SLA) it establishes with its customers. The SLA includes all contracted aspects of a service: financial, technical, operational, etceteras. Focusing on the technical aspect, a contract contains a customer-oriented

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: 10.1007/978-0-387-35620-4_43
D. Gaïti et al. (eds.), *Network Control and Engineering for QoS, Security and Mobility*

description of a service. For the purpose of automation, this technical description is formally contained in a standardised document named service level specification (SLS). Examples of SLSs are proposed in [2] and [3]. At the lower-end of the business layer, one or several SLSs are produced formally describing the customers' needs.

In order to implement the customer oriented SLS by provisioning it into the network, the SP must be able to translate it into an SP-oriented description. An approach to directly translate an SLS into configuration commands of each network element of the SP's network proves to be quite difficult. An interesting solution is provided by Policy-Based Management (PBM). The reason is not so much the usage of policy rules, but the high-level management view generally represented in existing policy information models (PIMs). The translation of an SLS into policy rules from appropriate policy information models is feasible as is illustrated in this paper. The resulting policy rules are SP-oriented. All SLSs are translated into policy rules at the upper-end of the network management layer.

Commonly these policy rules are used to populate a Policy Information Base (PIB). The PIB is a representation of abstract network functions (e.g. the dropping function of the DiffServ PIB [8]) based on a (vendor independent) network element (NE). As such, a PIB implements a policy information model partially or completely. The final configuration information of a NE is simply a specialisation of the configuration information contained in the PIB (associated with that NE). This completes the network implementation of a service that has started at the business layer as an SLS.

This paper describes all these steps in detail in the scope of an BGP/MPLS based IP VPN service. We present a textual description of an IP VPN service and its formal description (SLSs). Then we describe briefly the targeted PIMs and their implementations (PIBs). As a last step, we translate these SLSs into policy rules. These rules are defined in the targeted PIMs. We call this process Service Network Mapping (SNM).

Please note that, at the network level, the scope is limited to the configuration of the edge routers of the providers' networks supporting BGP/MPLS VPN [9] and MPLS/DiffServ for the transport of packets: we assume that the core of these networks is already configured. In addition, by stopping at the PIB level we remain vendor-independent.

2. THE IP VPN SERVICE

2.1 Description of the IP VPN

The IP VPN service considered hereafter is currently called "Provider Provisioned VPN" (PPVPN): a Service Provider (SP) provides his customer with a routed infrastructure that acts as a private IP cloud realised over a public IP infrastructure owned by this SP or other SPs.

To provision a PPVPN service the SP must be informed about the characteristics of the customers' VPN: its topology and reachability of VPN destinations. As the customer often uses private addresses on the LAN, the SP has to enable Network and port Address Translation (NAT) functionality to provide for routing connectivity on the public IP network. An SP may provide additional services such as firewall and encryption and offer SLAs on traffic flows (i.e. throughput and QoS) in the scope of site-to-site VPNs. An important issue is the granularity of QoS: the SP may offer aggregate SLAs or propose treatment of customer traffic on micro-flow level. Here a mixed solution is proposed: customer traffic is described in term of micro-flows within the SLS but a SP may aggregate this traffic once entering its core network.

To summarise, the PPVPN service requires the following points to be characterised: Reachability, NAT, Firewall, Encryption, and the definition of the transport of packets: VPN topology and QoS and bandwidth contract per link including traffic identification.

2.2 SLS of an IPVPN service

The SLS of a PPVPN service consists of a technical description of a generic packets transport service with QoS indicators and of a VPN specific part built on top of this transport service. In the following paragraphs we propose a description of the main components of an SLS formalising a generic transport service [2]. These components are:

- Service scope: a set of ingress and egress network interfaces of network nodes located at the edge of the provider's network and their relation to each other expressed as pipe, hose or funnel topologies. In this paper, these ingress and egress interfaces are also named input and output Service Access Points (SAPs).
- Flow description: identification of the traffic to be transported.
- Traffic profile and conformance test algorithm: for example token bucket.

- Excess treatment: dropping, shaping or remarking out-of profile traffic or ignoring the traffic.
- Performance guarantees: QoS defined as delay, jitter, packet loss, throughput.
- Service schedule: (de)activation time in terms of time of day, day of month, month of year, and any repetition that is needed.
- Reliability: maximum allowed downtime per TimePeriod and maximum time to repair.

In addition to these components and for addressing the PPVPN service, this SLS is to be completed with the following components:

- Reachability: defines the visibility of each site from each other site pertaining to the VPN.
- Firewall: defines what traffic is allowed or denied (a per-micro flow basis).
- NAT: network and port address translation in a per-micro flow basis.
- Encryption: defines encryption in a per-traffic basis.

Note that the description of a VPN service may require several (related) SLSs.

3. POLICY INFORMATION MODELS

Policy Information Models (PIMs) use policy rules for representing high level management functions. An important point is the abstraction of the representation used in a PIM. The goal is to let a network management operator easily define simple yet powerful rules to specify behaviour of network services. We use PIMs because it simplifies the mapping of the SLS to these policy information models.

A PIM is typically implemented as Policy Information Bases (PIBs). A PIB represents abstract network functions (e.g. QoS) of an abstract network element. However, for configuring a network, a PIB may not be needed in certain cases (e.g. using signalling to set up configuration).

It is important to define PIMs that are suitable for the implementation of a VPN service. For configuration of the edge of an MPLS/DiffServ network, the targeted PIMs are: an IPVPN PIM [11] strictly for the VPN part, QPIM [4] for the QoS part and PCIM/PCIMe [5][6] as background models. A description of these models and implementation is given in the next sections.

3.1 The PCIM and PCIMe models

PCIM and PCIMe introduce the basis for policy-based management: policy rule, condition, action, association, and group are defined in this

model. PCIM and PCIMe define in fact a meta-model and other PIMs, for example QPIM, specialise this model.

A common implementation is described in [7].

3.2 The QPIM model

QPIM defines policy rules for QoS in DiffServ or IntServ networks concerning the edge and the core of these networks. For DiffServ networks, QPIM defines traffic flow conditioning (i.e. testing and limiting egress traffic flows) at the edge, and per hop forwarding behaviour (bandwidth, queuing and dropping) of each router in the core.

QPIM is implemented as a PIB [8].

3.3 The IPVPN PIM

The IP VPN policy information model [11] is an Alcatel proprietary model anticipating the evolving IETF standard [12]. This model defines rules related to Reachability, NAT, Firewall and Encryption. It addresses only specific PPVPNs as specified in RFC 2547 bis [9] and is based on BGP/MPLS PPVPN.

This model is implemented as a PIB for a vendor independent representation of a BGP/MPLS VPN function of an NE called Provider Edge (PE) and is given in [13]. It has to be noted that the NAT, Firewall and security functions as discussed in the Alcatel PIM extensions are not included yet.

4. MAPPING THE SLS TO POLICY RULES

In the previous sections we have defined the relevant parts of the SLS, the targeted PIMs involved in defining the IP VPN service and a brief overview of the service implementation in the SP's network (edge only). This section fills the gap between the SLS and the policy models. It describes how the SLS is mapped, in a per-component basis, to the targeted PIMs.

The mapping of an SLS results in several policy rules. There is a direct relation between an SLS and the set of policy rules resulting from the mapping: modifying an SLS means updating this policy rule set, removing or adding one or several policy rules from this set or even a combination of the three.

The enforcement points must be defined for each policy rule in order to be able to enforce the policy at the network level.

Finally, the mapping is defined within the scope of the management of the edge: encompassing traffic conditioning, VPN Routing and Forwarding (VRF) functions. The mapping of the SLS components is presented according to these functions starting with traffic conditioning and then related to the VPN. But first we will examine the service schedule and the QoS components of the VPN service.

4.1 Service schedule component

A VPN service may be scheduled by defining a schedule component that is mapped as a specific simple condition from PCIMe: PolicyTimePeriod-Condition. This condition is only used with the mapping of other SLS components that are described in the following subsections.

4.2 QoS component

The QoS or performance guarantee component in a DiffServ network is simply represented by a DiffServ CodePoint (DSCP) [10]. Which DSCP to apply? This depends on the core network configuration, which is assumed to be already adequately pre-configured. Each DSCP is associated with a performance guarantee. A match between the performance guarantee defined in the SLS and a configured performance guarantee will give the DSCP value. The performance guarantees to be selected by the customer in his SLA and mapped to the SLS is directly mapped on the available core network performance guarantees (DSCP values). This DSCP value is used for defining the traffic conditioning (see next section). If there is no explicit performance guarantee, the default DSCP is set to a value corresponding to a best effort (BE) service.

4.3 Traffic conditioning related components

Traffic conditioning is defined according to the scope, flow identification, traffic profile, and excess treatment components of the SLS and requires the DSCP resulting from the mapping of the performance guarantee component defined in the previous section.

Traffic conditioning is implemented as a single policy rule (named TrafficConditioning in *Figure 1*). The compound condition of this rule implements the flow identification and schedule components. The action part implements the Traffic Profile and the conformance algorithm component.

The resulting policy rule is to be enforced on each interface identified as SAP in the scope component of the SLS. That is, this policy rule is to be

enforced on the traffic entering if an interface is identified as an ingress SAP and on the outgoing traffic if an interface is identified as an egress SAP.

In terms of the models, traffic conditioning is implemented as an instance of the class PolicyRule [5].This policy rule is built as follows:

- The condition part of this rule is an instance of the class Compound-PolicyCondition [6] associated with the conditions on the flow identification and on the validity period by means of an instance of the class PolicyConditionInPolicyCondition [6]. The flow identification is mapped to at least one simple condition: a simple condition is defined by using a PolicyVariable [6] and a PolicyValue [6]. The validity period condition is defined in the "Schedule" subsection.
- The action part of this rule is an instance of the class QosPolicyPolice-Action [4] implementing the conformance test to a traffic profile and the associated actions to trigger. Possible actions on out-of-profile traffic are transmitting, dropping, remarking, or shaping. A possible action on in-profile traffic is marking with a given DSCP computed as specified in the subsection named "QoS component". *Figure 1* illustrates the generic form of a policy rule implementing traffic conditioning.

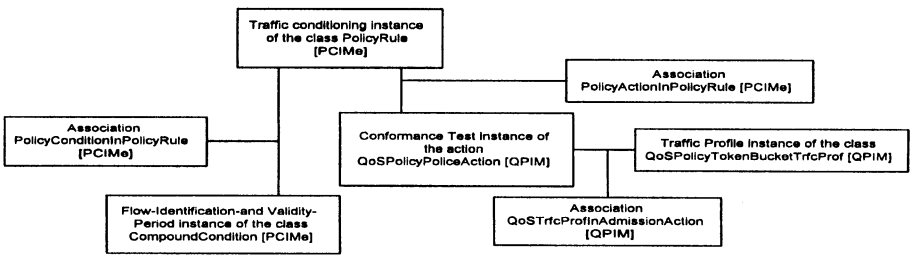

Figure 1. Traffic conditioning policy rule

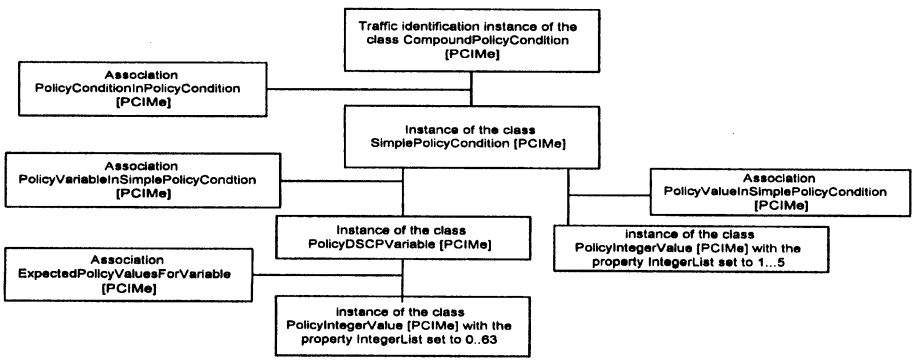

Figure 2. Flow identification mapping example

Figure 2 illustrates an example of mapping of a flow identification component describing an ingress traffic of packets marked with a DSCP value between 1 and 5.

Finally, *Figure 3* is an illustration of the implementation of the action performed on in-profile traffic.

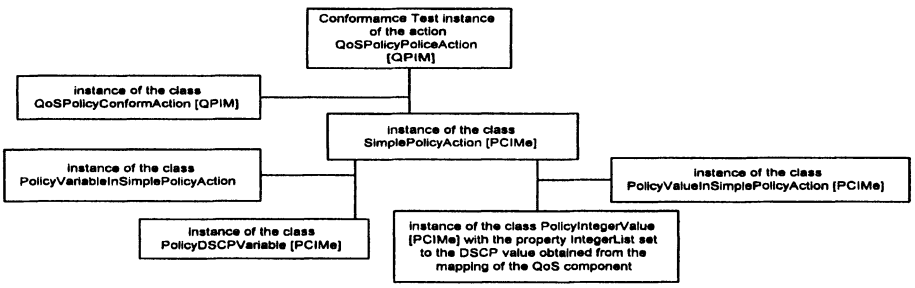

Figure 3. Marking operation of in-profile traffic

4.4 VPN related components

The mapping to the IP VPN model may result in several policy rules defining reachability, NAT, Firewall and Encryption. These are the main rules that have to be enforced on the interface corresponding to the SAP defined in the scope component of the SLS. In contrast, policy rules defining reachability are enforced on the appropriate PE by the respective VRFs.

4.4.1 Reachability

The reachability component can be implemented using only one or using several policy rules: one rule per VRF. In any case, each rule is defined similarly: the condition part of this rule is a time validity period and the action part is one action, ProvisionVRFPolicyAction, as defined in [11]. The only property of this action is the list of SAPs (to associate with the VRF) as defined in the reachability component.

The rule we created is to be enforced on the appropriate PEs that are deduced from the definition of the reachability component of the SLS. *Figure 4* gives an illustration of this rule.

The Reachability component also defines the routing of packets within a VPN and can again be implemented as one or several policy rules. Each rule represents the accessibility of source sites to destination sites (reachability rules are uni-directional). Thus, the number of required rules depends on the definition of the accessibility. Each rule illustrated below is defined similarly: the condition is the time validity period and the action is one

action, ConfigureVRFPolicyAction, as defined in [11]. This rule is to be enforced on PEs attached to source sites.

Figure 5 is an illustration of a rule that allows the distribution of routes from site A to site B only.

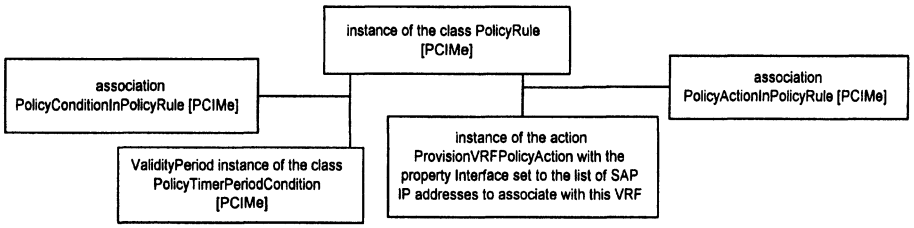

Figure 4. A VRF creation

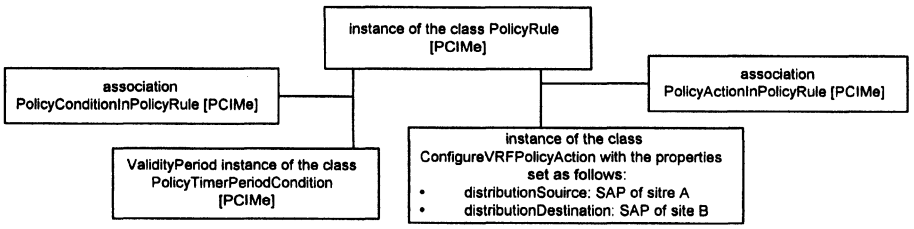

Figure 5. A route distribution

4.4.2 Network and Port Address Translation

The NAT component is implemented as a single policy rule that is defined as follows. The rule condition part is the result of the mapping of the flow identification and schedule components (see traffic conditioning subsection). The action part is a CompoundPolicyCondition [6] of NATPolicyAction. The only two properties of the NATPolicyAction are set as follows:

- translateFromIPV4Address property is set to the value of an IP address to translate from
- translateToIPV4Address property is set to the value of a global IP address to translate to

Each resulting policy rule is to be enforced on the interface(s) of the edge node(s) implementing SAP(s) identified in the scope component of the SLS. *Figure 6* gives an example of a rule translating a local IP address to a global IP address. This rule is to be enforced on traffic identified in the condition and during a given validity period.

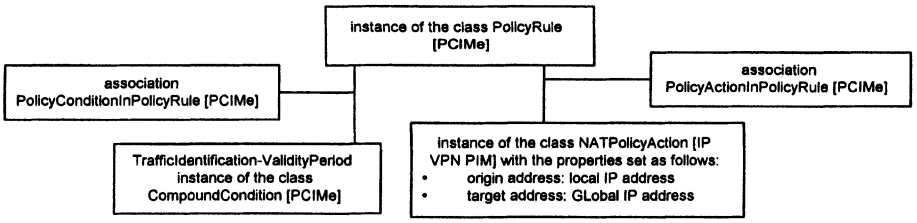

Figure 6. Network and Port address translation mapping example

4.4.3 Firewall

The firewall component maps to one policy rule as follows. The condition part of this rule is the result of mapping the flow identification and schedule components (as defined for the traffic conditioning). The action part is a FirewallPolicyAction. This policy rule is to be enforced on concerned PE(s) providing input SAP(s) defined in the scope component.

Figure 7 is an illustration of a rule denying (with notification) an identified traffic during a given validity period.

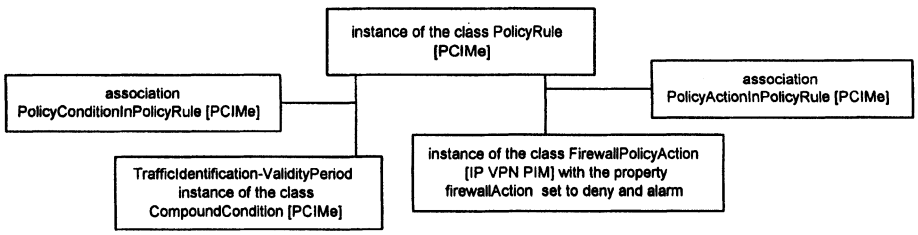

Figure 7. Firewall component mapping example

4.4.4 Encryption

Encryption also results in one policy rule of which the condition part is the result of the mapping of the flow identification and schedule components and the action part is an instance of the action EncryptionPolicyAction. This policy rule is to be enforced on input SAPs defined in the scope component of the SLS. Note that the output SAPs are used in the rule definition.

Figure 8 is an illustration of a rule encrypting identified traffic using a DES algorithm during some validity period.

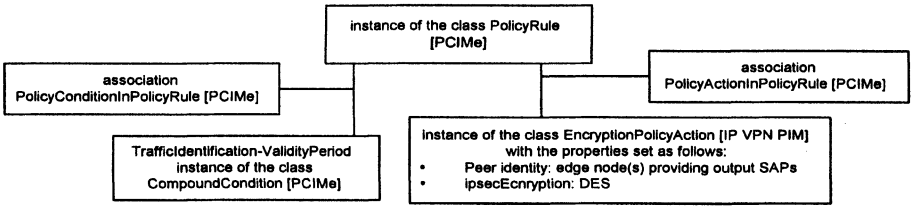

Figure 8. Encryption component mapping example

5. MAPPING FROM PIB DATA TO DEVICE CONFIGURATION

The network device abstracted Policy Rules in the PIB need to be translated into specific network element (NE) configuration commands. This mapping is straightforward if the NE is 'policy aware' and supports a PIB natively. This NE will (partially) retain a local copy of the PIB and no translation is necessary, only a distribution to the device PIB.

In the case of a legacy NE that is not 'policy aware', still the majority of existing NE's, a translation of the PIB information is necessary at the distribution point into SNMP messages or other formats and protocols for configuration of the NE. In our research we have validated the scenario of configuring a legacy router using mapping of policy into "Command Line Interface" commands.

6. CONCLUSION

This document demonstrates that it is possible to completely implement a service from a formal Service Level Specification down to the configuration of network elements. The following points can be concluded:

a) The SLS template is very powerful for formally expressing customers requirements. Independent of network implementation details.
b) Policy-Based Management is an adequate and efficient approach for specifying a high level configuration of the provider network. The gain in efficiency is achieved through the ability to simply implement an SLS by means of policy rules and subsequent enforcement of these rules on the network without technological consideration. This requires careful design of the policy information models to support service design and implementation. These models are PCIM [5], PCIMe [6], QPIM [4], and the IP VPN PIM [11].

Another result is the validation of the use of policy information model extensions to the basic set of models defined by the IETF for configuring additional service requirements in the network.

What comes next? Our IP VPN model can evolve to a more generic model covering other possible implementations, e.g. tunnel-based VPNs. We also have to take into account the reliability and assurance aspects of a VPN service. Another important step is to extend the scope of the information models to the core network. The work to be done is similar to the work realised now: specifying an appropriate model of the core configuration and implementing this model, which will likely be related to a bandwidth broker and an admission control function.

Finally, this paper illustrates the work done within Alcatel on flow through service provisioning in the network directly using the customer needs as input. This is an important step towards a customer-oriented and perhaps customer driven management solution.

7. REFERENCES

[1] TeleManagement Forum, GB 910, "Telecom Operations Map", March 2000

[2] IST-1999-11253-TEQUILA, Traffic Engineering for QoS in the Internet , at Large Scale, http://www.ist-tequila.org

[3] EURESCOM-P1008, Inter-operator interfaces for ensuring end to end QoS, http://www.eurescom.de/public/projects/P1000-series/p1008

[4] IETF-draft, "Policy QoS Information Model", draft-ietf-policy-qos-info-model-04.txt, Y. Snir, Y. Ramberg, J. Strassner, R. Cohen, B. Moore, November 2001, work in progress

[5] IETF-RFC 3060, "Policy Core Information Model -- Version 1 Specification", B. Moore, E. Ellesson, J. Strassner, A. Westerinen, February 2001

[6] IETF-draft, "Policy Core Information Model Extensions", draft-ietf-policy-pcim-ext-08.txt, B. Moore, May 2002, work in progress

[7] IETF-RFC 2753, "A Framework for Policy-based Admission Control", R.Yavatkar, D. Pendarakis, R. Guerin, January 2000

[8] IETF-draft "Differentiated Services Quality of Service Policy Information Base", draft-ietf-diffserv-pib-06.txt, M. Fine, K. McCloghrie, J. Seligson, K. Chan, S. Hahn, C. Bell, A. Smith F. Reichmeyer, March 2002, work in progress

[9] BGP/MPLS VPNs, draft-ietf-ppvpn-rfc2547bis-01, E.C. Rosen et. al., July 2002, work in progress

[10] IETF-RFC 2475 "An Architecture for Differentiated Services", S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, December 1998, informational

[11] NET-CON'2002, "A Policy Information Model for RFC2547-like IP VPNs", A. Gonguet, O. Poupel, October 2002

[12] IETF-draft, "IP VPN information model", draft-iyer-ipvpn-information-model-01, M. Iyer, A. Gonguet, C. Jaquenet, P. lago, R. Scandarioto, February 2002, work in progress

[13] IETF-draft, "BGP/MPLS VPN Policy Information Base", draft-yacine-ppvpn-2547bis-pib-00.txt, Yacine El Mghazli, April 2002, work in progress