

# ExSERT: Enabling Distributed Monitoring at Internet Exchange Points

Felipe Huici Saleem N. Bhatti<sup>†</sup> and John Souter<sup>‡</sup>

<sup>†</sup>University College London <sup>‡</sup>London Internet Exchange

**Abstract:** Internet Exchange Points (IXPs) carry a significant portion of Internet traffic. One of their principal activities is to monitor their networks to ensure their effective and reliable operation. In particular, Euro-IX, an association of most European IXPs, has a requirement to perform distributed monitoring and to be able to share monitoring tools and data among its members. However, currently available tools are either commercial point monitoring tools or do not provide a mechanism for easy creation and sharing of new monitoring tools and data. In this paper we present ExSERT, the *Extensible Secure Event and Report Toolkit*, which will provide solutions to these problems.

## 1 Introduction

An Internet Exchange Point (IXP) is a physical infrastructure that allows different Internet Service Providers (ISPs) to exchange traffic between their autonomous systems by means of mutual peering agreements. Any ISP that is connected to the IXP can exchange traffic with any of the other ISPs connected to the IXP; this is done using a physical connection to the IXP, thus overcoming the scalability problem of having multiple individual interconnections between ISPs. Also, by enabling traffic to take a more direct route between many ISP networks, an IXP can improve the efficiency of the Internet, resulting in a potentially better service for the end user. Furthermore, since many networks have more than one connection to the Internet, it is not unusual to find several routes to the same network available at an IXP, thus providing a certain amount of fault tolerance.

Since it acts as a central point for traffic exchange, an IXP must ensure the effective and reliable operation of its network infrastructure. IXPs are typically quite small businesses, with perhaps surprisingly few staff given the critical part they play in the infrastructure of the Internet. This places heavy pressure on them to use monitoring tools and automation, since they rarely if ever have the resources to run a 24/7/365 network operations center. IXPs therefore rely on software tools which allow them to monitor network traffic as part of a wider network management system and inform them of any occurring or potential problems. To ease the troubleshooting process it would be highly beneficial if IXPs could exchange information and monitor other IXPs' traffic to share knowledge of past problems as well as current traffic patterns. Access to data or data interfaces is very jealously guarded, so one of the paramount requirements of any mechanism that facilitates such an exchange is that the transfer happens in a secure and confidential manner, so that only the intended parties have access to the data. Euro-IX (the European Internet Exchange Association), an association of most of the IXPs in Europe, provides a perfect forum for such an exchange of data among its members. Our aim is to enable distributed monitoring at IXPs by implementing ExSERT, an extensible, secure monitoring toolkit based on a flexible, modular architecture.

## 2 Motivation

One of the most important aspects of the day-to-day operations of an IXP is monitoring its network to guarantee that it is running as expected. An important objective of the Euro-IX association is to be able to share not only tools developed by IXPs but also data collected from monitoring tools. However, because of their commercial nature, the available distributed monitoring tools do not inter-operate, requiring that all IXPs adhere to one tool in order to achieve the objective of data sharing; this level of homogeneity is clearly impractical. To make matters worse, commercial software is difficult to tailor to specific needs: the user has to rely on the given APIs which may or may not be flexible enough. For these reasons, a non-commercial toolkit is required.

Another problem with the current state of affairs is that even though Euro-IX members are willing to share tools developed in-house, these are typically tailored to an IXP's specific hardware, and, consequently, cannot be used by another IXP without re-writing a significant portion of the code. Many of the tools

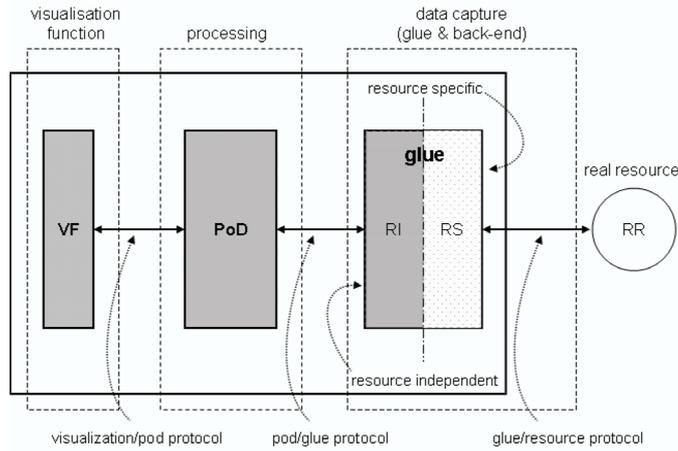


Figure 1: The Remote Monitoring Function (RMF).

that members of Euro-IX have developed and use on a daily basis are semantically similar and differ, for the most part, only syntactically. A mechanism is needed, then, that will take advantage of these semantic similarities in order to provide easy sharing of tools among members.

Yet another incentive for enabling distributed monitoring at IXPs is that it would significantly benefit research into areas that, by their very nature, depend on the gathering of distributed data. Examples include the interaction between BGP convergence and route-flap damping; the large-scale effects of interaction between inter-domain and intra-domain routing; the detection of distributed denial-of-service attacks as well as determining their effects on the network; and congestion control and traffic engineering in the large. Further, in [1] the authors mention the lack of monitoring tools that will measure *aggregate* behavior of nodes in the system; ExSERT would fill this gap.

While a large number of network monitoring tools exist ([2][6]), none of them resolve all of the conflicts described. A study by the London Internet Exchange (LINX), one of the largest IXPs in Europe, concluded that no monitoring tool existed that met their requirements regarding sharing of data, easy sharing and creation of tools and distributed monitoring. Although the study was informal and internal, it was extensive, covering not only similar IXPs throughout Europe, but also including LINX's membership of ISPs and telecommunication companies. Although LINX identified some reasonably comprehensive commercial tools (e.g. SMARTS [7]), these were very expensive and thus out of reach for most IXPs, and also seemed to require a very high level of commitment and data conversion, with no absolute assurance that they would fit well to IXP requirements. Generalizing, LINX found that most IXPs and ISPs used a mixture of open-source tools such as MRTG [3], RRD [5], SNIPS [8] and Nagios [4], with a great deal of local customization and local integration effort. Clearly, then, a new flexible and easily extensible toolkit is needed.

### 3 ExSERT

In order to address the issues discussed in the previous section, we propose ExSERT, the *Extensible Secure Event and Report Toolkit*. The toolkit provides a simple and powerful way of creating new tools that will report data in a fast, reliable and secure way. In subsection 3.1 we discuss the details of this architecture and in subsection 3.2 we describe the different components of the toolkit itself. The implementation of the toolkit is currently under way.

#### 3.1 Architecture

The system's modular architecture revolves around a Remote Monitoring Function (RMF), an abstraction that performs point monitoring of a particular aspect of a network, such as round-trip time measurements. Each RMF contains three elements, a glue, a PoD, and a visualization function (VF), all of which communicate using well-defined protocols (see figure 1). This modularity yields great flexibility, allowing each of these components to physically reside on separate pieces of hardware and on different networks.

The glue consists, in turn, of two components. The resource-specific element communicates with the real

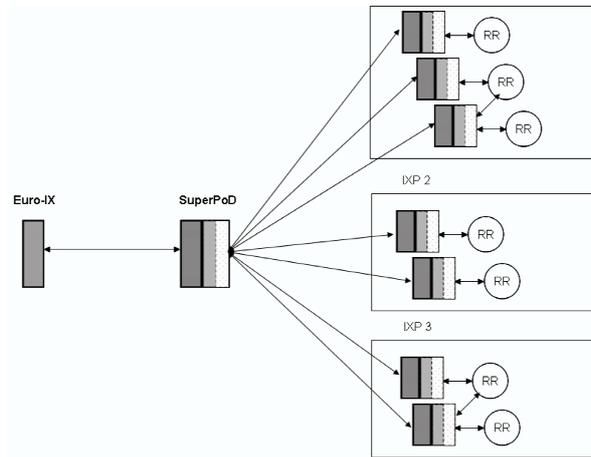


Figure 2: Aggregating similar PoDs into a superPoD.

resource (which could be a piece of hardware, a log file or even a measuring tool such as ping). The resource-independent element then adapts the data from their native format to a well-defined format. The split between a resource-specific element and a resource-independent one yields easy portability: if an IXP develops a tool that it wants to share with another IXP, the other IXP need only rewrite the resource-specific part of the glue to match its hardware; all the other components of the RMF remain the same.

The PoD performs two functions: it analyzes the data supplied by the glue and acts as a light-weight server that provides the analyzed data to clients (the analysis could be a simple function such as averaging a set of round-trip time measurements). Finally, the visualization function connects to the PoD, retrieves the data, and displays it, perhaps graphically.

While the description so far has consisted of only one RMF, the architecture's modularity permits much more involved monitoring, including aggregating similar PoDs so that their output becomes the input to the glue of a new element called a superPoD. For instance, all Euro-IX members could run a PoD that measures aggregate throughput at that IXP; a superPoD running at Euro-IX headquarters could then periodically connect to all of these PoDs and add up the throughputs to provide an overall Euro-IX traffic rate (see Figure 2). Incidentally, this gathering of throughputs is currently performed at Euro-IX, though the process is done in an ad-hoc manner and is very unreliable and inconsistent despite its being quite useful to its members.

### 3.2 The Toolkit

While it would be possible to implement a new glue, PoD and VF for each tool or RMF desired, a considerable portion of the code is common to all RMFs, and ideally the process of creating new tools should be automated as much as possible; this is precisely the goal of ExSERT and its compiler.

A network operator wishing to develop a new distributed monitoring tool uses an XML schema to specify the type of data that the monitoring reports will contain, based on some basic ExSERT schema types. In addition, the operator specifies a few tool-specific configuration parameters, such as what port the PoD should be listening on and what the target source code language should be. This information is then used as input to the ExSERT compiler, which creates all the necessary source code (see figure 3). The operator only needs to write the resource-specific part of the glue, potentially allowing for the creation of new tools in a matter of minutes. In addition, if the generated tool did not precisely match the requirements of the operator, he or she could modify since the source code is available. Further, researchers with access to IXPs could gather substantial amounts of distributed, real-world data having spent a minimal amount of time creating the actual tool.

### 4 Future Work

The proof-of-concept implementation in Java of the architecture has been running stably and taking round-trip time measurements at LINX since August 2004. It has been tested with a client that gathered

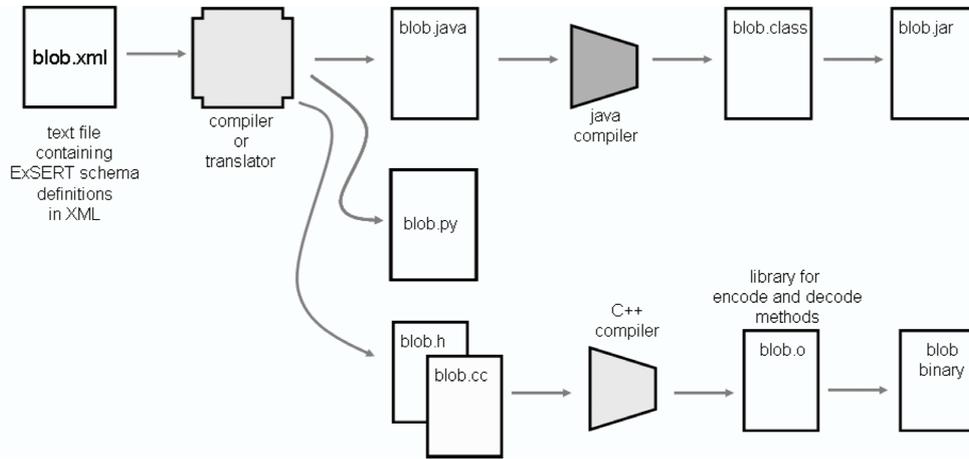


Figure 3: The ExSERT compiler.

data from LINX and from a dummy IXP set up at University College London and displayed them graphically. Future steps include a complete rewrite of the code in Python, this time including the implementation of ExSERT's tool-creation capabilities and compiler. This new implementation will also focus on scalability, performance and on designing a better security model. We have contacts with several members of Euro-IX that expressed their interest in ExSERT when it becomes available, so real-world deployment is planned.

### Acknowledgments

Agron Fazliu, Felipe Huici, Alexander Papitsch, Andreas Protopapas and Giorgos Savvides designed the architecture and carried out the proof-of-concept implementation as part of their master's thesis.

### References

- [1] Ankur Jain, Joseph M. Hellerstein, Sylvia Ratnasamy, and David Wetherall. A wakeup call for internet monitoring systems: The case for distributed triggers. In *Proc. 3rd ACM SIGCOMM Workshop on Hot Topics in Networks (HotNets)*, San Diego, CA, November 2004.
- [2] MonitorTools.com. Network monitor software and windows development tool (web site). <http://www.monitortools.com>.
- [3] MRTG. The Multi Router Traffic Grapher. <http://people.ee.ethz.ch/~oetiker/webtools/mrtg/> (Web site).
- [4] Nagios. Nagios Home (Web site). <http://www.nagios.org/>.
- [5] RRD. RRDtool - About RRDtool (Web site). <http://people.ee.ethz.ch/~oetiker/webtools/rrdtool/>.
- [6] Stanford Linear Accelerator Center Network Monitoring (Web site). Network Monitoring Tools (Web site). <http://www.slac.stanford.edu/xorg/nmtf/nmtf-tools.html>.
- [7] SMARTS. SMARTS - The Leading Provider of Automated Business Assurance Solution (Web site). <http://www.smarts.com>.
- [8] SNIPS. SNIPS - System and Network Management Monitoring Software (Web site). <http://www.netplex-tech.com/snips/>.