# Network Anomaly Detection in the Cloud: The Challenges of Virtual Service Migration

Kirila Adamova, Dominik Schatzmann, and Bernhard Plattner
ETH Zürich
Zürich, Switzerland
Email: kirila.adamova@gmail.com
{dominik.schatzmann, plattner}@tik.ee.ethz.ch

Paul Smith
AIT Austrian Institute of Technology
2444 Seibersdorf, Austria
Email: paul.smith@ait.ac.at

***Abstract*—The use of virtualisation technology in the cloud enables services to migrate within and across geographically diverse data centres, e.g., to enable load balancing and fault tolerance. An important part of securing cloud services is being able to detect anomalous behaviour, caused by attacks, that is evident in network traffic. However, it is not clear whether virtual service migration adversely affects the performance of contemporary network-based anomaly detection approaches. In this paper, we explore this issue, and show that wide-area virtual service migration can adversely affect state of the art approaches to network flow-based anomaly detection techniques, potentially rendering them unusable.**

## I. Introduction

Cloud computing has proved to be a popular way for organisations to provision services for their users. There are a number of reasons for this popularity, including potential reductions in operating costs, flexible and on-demand service provisioning, and increased fault-tolerance. Drawn to these benefits, operators of critical information infrastructures – the ICT infrastructures that support gas and electricity utilities and government services, for example – are considering using the cloud to provision their *high assurance* services. This is reflected in a recent white paper produced by the European Network and Information Security Agency (ENISA), which provides specific guidelines in this area [1].

Deploying high assurance services in the cloud increases cyber-security concerns – successful attacks could lead to outages of key services that our society depends on, and disclosure of sensitive personal information, for example. To address these concerns, a range of security measures must be put in place, such as cryptographic storage and network firewalls. An important measure is the ability to detect when a cloud infrastructure, and the services it hosts, is under attack via the network, e.g., from a Distributed Denial of Service (DDoS) attack. A number of approaches to network attack detection exist, based on the detection of *anomalies* in relation to normal network behaviour [2].

One of the essential characteristics of cloud computing is the use of virtualisation technology, which supports the migration of services across a physical infrastructure within and between large-scale cloud data centres – known as *local* and *wide-area* migration, respectively. The reasons for service migration are manifold, including responding to hardware faults, planned maintenance tasks, and handling localised peaks in service requests by moving services "closer" to their user. Whilst virtual service migration has a number of benefits, it has the potential to make the implementation of security measures challenging, therefore introducing new vulnerabilities [3].

In this paper, we are specifically interested in examining the affect virtual service migration has on contemporary network anomaly-based attack detection techniques – as services move, migration may be observable in the network traffic that is being used for anomaly detection. As mentioned earlier, such techniques aim to detect anomalous traffic in relation to a learned baseline, which represents normal behaviour. It is unclear to what extent virtual service migration, which is arguably representative of "normal" cloud behaviour, can be incorrectly observed as an anomaly, and therefore an attack. Conversely, attacks may be missed because of virtual service migration. If this problem is significant, such detection techniques could be rendered unusable for the cloud, thus representing a significant vulnerability and a potential inhibitor to the deployment of high assurance services.

Using a novel toolchain, which simulates attacks and virtual service migration in network flow traces, we have examined the detection performance of two anomaly detection techniques – Principal Component Analysis (PCA) [4], [5], [6] and the Expectation-Maximisation (EM) clustering algorithm [7], [8]. In previous research, these detection techniques have been shown to give acceptable detection performance results in non-cloud settings. Under different attack and virtual service migration scenarios, we have measured their ability to reliably detect attack behaviour in the cloud. Our results suggest that, in some configurations, a potentially insecure number of attacks are missed, and an unusably high number of alarms pertaining to normal behaviour are generated. This result draws into question the use of these techniques, and potentially others, in large cloud data centres, in which virtual service migration is a common undertaking.

The rest of this paper is organised as follows: Section II discusses related work – our investigations indicate that, to the best of our knowledge, there is no previous work that directly addresses the problem explored in this paper. A discussion on virtual service migration and its affect from a network perspective is presented in Section III. Section IV describes the toolchain and traffic data that we used to obtain the experimental results, which are described in Section V. We conclude and discuss potential solutions to the problem we have explored in Section VI.

## II. Related Work

A number of approaches can be used to detecting network-borne attacks, such as Distributed Denial of Service (DDoS) attacks, to large cloud data centres. There has been significant research interest in and deployment of algorithms that aim to identify deviations from normal traffic behaviour – *anomalies* – that are indicative of attack behaviour. A survey of anomaly detection approaches has been produced by Chandola *et al.* [2]. Our research is based on contemporary approaches that use spectral analysis to detect anomalies in entropy measures, derived from network flow summary data. Tellenbach *et al.* examine the use of Kalman filter, Principal Component Analysis (PCA), and Karhunen-Loève Expansion (KLE) when considering the affect of different entropy measures on detection performance [4]. For our investigation, we use PCA, as it has been shown to perform well when configured appropriately [6] and continues to be investigated in the research community [5]. To explore the potential extent of the problem that we have identified, we have used an anomaly detection algorithm that uses clustering; specifically, the Expectation-Maximisation (EM) algorithm, which has shown to give promising detection performance [7], [8].

Our investigation relates to identifying potential shortcomings in flow-based anomaly detection techniques, which manifest due to their deployment; in this case, wide-area virtual service migration in large cloud data centres. This line of enquiry is, in-part, motivated by previous research, conducted by Brauckhoff *et al.*, which examined the impact sampling of network flow data has on anomaly detection [9]. Their study showed that statistical techniques that identify anomalies in traffic volumes perform less effectively under sampling conditions. Furthermore, they suggest spectral-based analysis, using entropy measures of traffic feature distributions, e.g., source and destination IP address and port numbers, are more robust to sampling. As we will discuss in Section III, wide-area virtual service migration manifests as a change in network traffic volume, observable at a data centre – this is similar to the affect sampling has. This observation was one of the motivations for the choice of a PCA-based approach for the study we present in Section V. To the best of our knowledge, our investigation is the first to examine the affect virtual service migration has on network flow-based anomaly detection techniques.

## III. Preliminaries – Virtual Service Migration

In order to consider the analysis results that are presented in Section V, a brief discussion on virtual service migration is required, along with its affect on data centre network traffic.

### A. Virtual service migration in large cloud data centres

An essential characteristic of cloud computing is the use of virtualisation technology [10], whereby services execute in virtual machines atop of a physical compute, storage and network infrastructure. Virtualisation supports the ability to migrate services between different underlying physical infrastructure. There are multiple reasons to migrate services as a consequence of the *day-to-day* operation of a large data centre, including load balancing, failure of the underlying hardware [11], in response to routine maintenance tasks, and reducing network costs. In some cases, a service must be migrated between geographically and topologically distinct data centres. For instance, if the majority of client connections for a service originate in Asia, but the service is hosted in Europe, network costs can be decreased if the service is moved topologically closer to its clients. Furthermore, geographical diversity of data centres is supported by commercial cloud providers, such as Amazon, in order to improve fault-tolerance. Moving services between data centres is known as *wide-area* migration. Conversely, *local-area* migration occurs when a service is migrated within a data centre. In both cases, there are multiple approaches to ensure the network traffic that is destined for a migrated service is forwarded to the correct location [12].

### B. Observing migration in network traffic

Importantly for anomaly detection, service migration may result in observable effects in network traffic – this depends on where network traffic that is to be analysed is collected in the data centre topology, and the type of migration that is carried out. For example, if a local-area migration is executed, the change in traffic could be observable at Top-of-Rack (ToR) and aggregate switches, but not at the gateway to the data centre. (This is the case if we assume a data centre topology, as outlined in[13].) Since it is common practice to analyse traffic at the edge of a data centre, local area migration is opaque at this location. However, for wide-area migration, traffic destined to the migrated service will be forwarded to a different data centre, and will stop being received at the origin after the migration process has finished. This will result in potentially observable effects in traffic collected at the edge of a data centre, and hamper anomaly detection techniques. In this paper, it is precisely this problem that we explore – i.e., the effect that wide-area service migration has on anomaly detection techniques when network traffic that is used for analysis is collected at the data centre edge.

Broadly speaking, what occurs at a data centre after a wide-area migration can be described as *removing* all the traffic related to a migrated service at the source data centre, and *adding* new traffic for a service or set of services at the target data centre. We acknowledge there will be other observable effects, such as virtual machine state being transferred, but we expect these to be relatively minor in a large data centre. These effects are illustrated in Fig. 1, which shows changing source port entropy values over time with a dataset that includes a service migration at time $10/24$ and a port scan anomaly injected. These plots were created using the dataset and toolchain we describe in Section IV. The migration can be seen as the entropy decreases at the source data centre (Fig. 1(a)), where a service is removed and the source port distribution is not as previously dispersed. At the target data centre (Fig. 1(b)), the inverse can be observed, including port entropy changes that could result in false alarms. Another interesting observation that can be made is how the injected anomalies appear after migration – in Fig. 1 these are shown as the peaks. At the source data centre, after migration, the anomalies appear more pronounced when compared to the rest of the traffic. The inverse can be observed at the target data centre after migration.

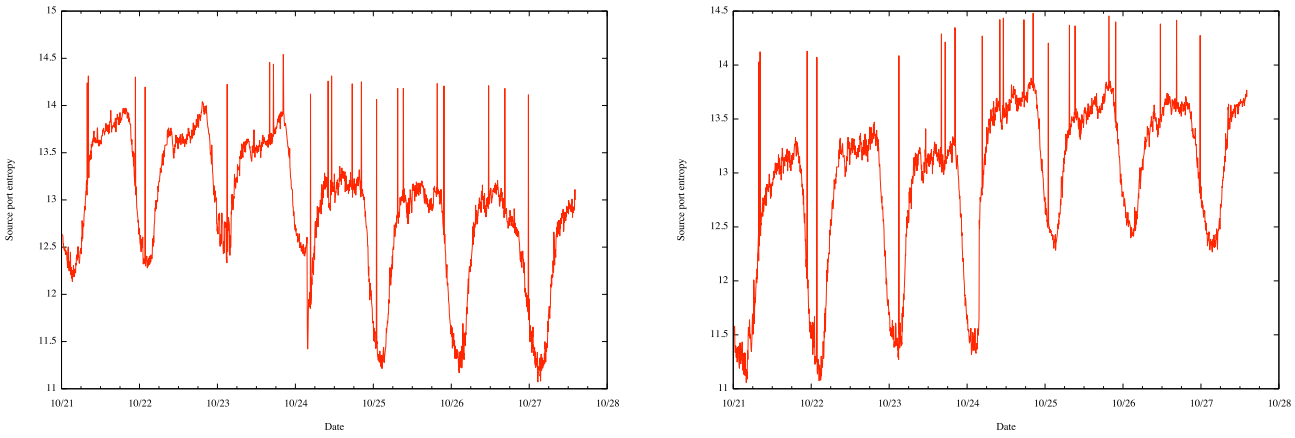


(a) Source data centre

(b) Target data centre

Fig. 1. The effect of migration on source port distribution of traffic at the source and target data centres. Migration occurs at $10/24$; the dataset includes a port scan.

## IV. Simulating Cloud Service Migration

In order to examine the affect virtual service migration has on anomaly detection, we created a toolchain that can be used to "simulate" virtual service migration and attacks in a given network flow dataset. The toolchain integrates a number of existing software to achieve this goal. It was necessary to develop the toolchain as we were *(a)* unable to find any public datasets that contained migration behaviour for a large data centre; and *(b)* it afforded us a degree of flexibility for our experiments, e.g., to inject migration at arbitrary times in a baseline dataset, with varying numbers of services being migrated. We briefly describe the dataset we used and our toolchain.

### A. Traffic dataset

For our experiments, we used network flow data that was collected from the Swiss national research network – SWITCH[1] – a medium-sized backbone network that provides Internet connectivity to several universities, research labs, and governmental institutions. The flow data was collected, without explicit sampling configured, at the border routers of the SWITCH network. We used a week's worth of data from 21–28 October 2012 for our experiments. To produce flow data that is representative of that seen at a large data centre, we extracted the top 200 TCP-based services, based on total number of flows, from the dataset. We confirmed these services represented those seen at a cloud data centre by comparing their features to a subset which accessed services at one of Amazon's European data centres. We found them to be comparable. Using this baseline dataset we can then include migration and attack behaviour using our toolchain.

### B. Experimentation toolchain and method

Fig. 2 presents an overview of the toolchain that we developed for our experiments. Initially, NetFlow data is processed using an extended version of a flow processing framework, called Flowbox[2]. This process includes filtering the top 200 selected services, and storing the flow data for each service in separate files that span five minute periods. Anomalies are generated using the FLAME tool [14], and stored in NetFlow format. We conducted experiments with a volume-based attack, i.e., a Distributed Denial of Service (DDoS) attack, and non-volume based vertical and horizontal port scans. Similarly, anomalies that are to be injected into the dataset are stored in five-minute files.

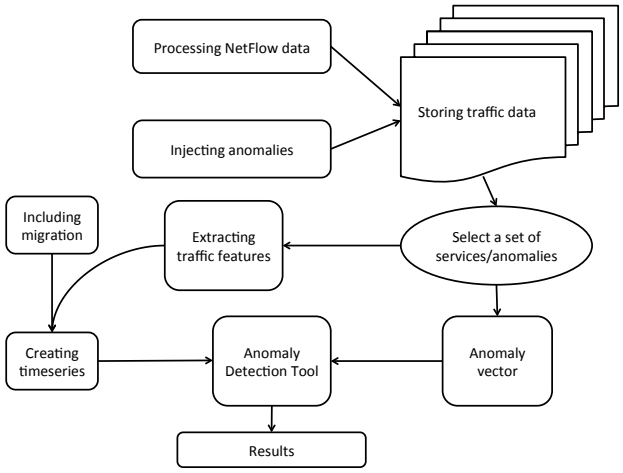


Fig. 2. Overview of the toolchain used for our experiments

The *anomaly detection tool* that analyses the traffic data, along with the injected anomalies and migration, has three inputs: an *anomaly vector*, which specifies at what times anomalies were injected, and two *time series* files – a baseline and evaluation file. A time series file contains entries that summarise each five minute period of the traffic. This description includes seven traffic features and a timestamp. These time series are created by extracting the traffic features from a set of pre-selected services and anomalies. The traffic features are *volume-based*, the flow, packet and byte count, and *distribution-based*, the Shannon entropy of source and destination IP and port distributions. These features are commonly used in flow-based anomaly detection.

[1] http://www.switch.ch/

[2] https://github.com/FlowBox

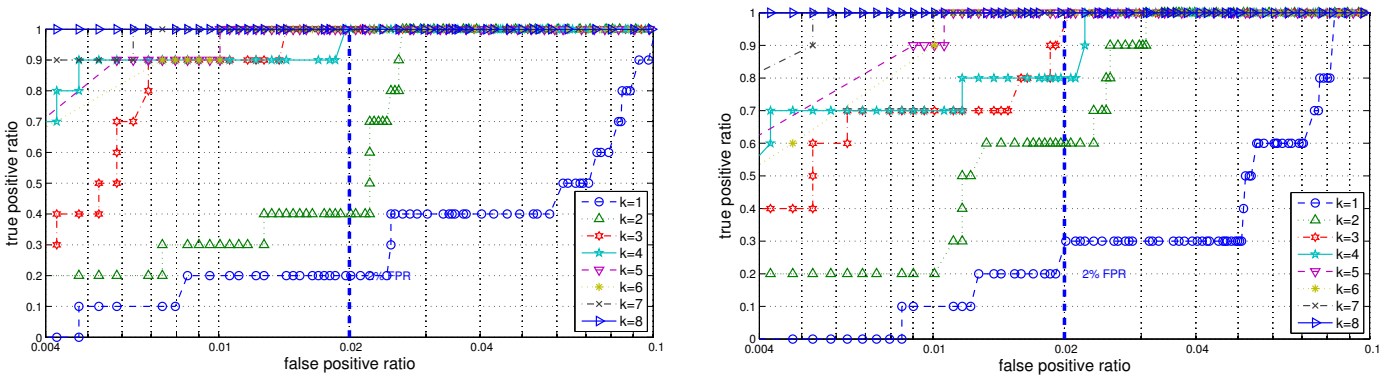


(a) ROC curve *without* virtual service migration

(b) ROC curve *include* virtual service migration

Fig. 3. ROC curves that show the performance of a Principal Component Analysis (PCA)-based anomaly detection approach to detect a Distributed Denial of Service (DDoS) attack, both without (3(a)) and with (3(b)) virtual service migration.

Time series entries for a five minute period are created by counting the number of flows, summing the number of packets and bytes, and creating map structures of IP address and port distributions. The latter are used to calculate an entropy measure. When all the files that describe services over a five-minute period have been processed, entropy is calculated and all features are stored as one time series entry. The aforementioned baseline time series are created for training data; these are free from anomalies and migration. Conversely, evaluation time series contain anomalies and may contain migration, depending on the experiments we wish to conduct.

As discussed in Section III, wide-area migration can be observed as the removal or addition of new services and associated inbound and outbound traffic at the source and target data centre, respectively. We simulate this behaviour as follows: at the source data centre a set of services is selected for the baseline, and a further set are marked for migration at a given time $t$. For the evaluation time series the creation starts without any changes. However, when $t$ is reached in the dataset, the time series are created without services that were marked for migration. This simulates the stopping of the migrated services at the source data centre. Meanwhile, for the target data centre, the process is the reverse. When $t$ is reached, the set of services on which the time series are based will further include the set of migrated services.

It can be seen that using this toolchain we can flexibly create evaluation scenarios that include a range of attack behaviours, using the FLAME tool, and service migration activity of different magnitude. These scenarios can then be provided as input to different flow-based anomaly detection techniques, as discussed in the following section.

## V. The Impact of Migration on Detection Performance

In this section, we discuss the affect wide-area virtual service migration has on two anomaly detection approaches: spectral analysis, based on Principal Component Analysis, and clustering using Expectation-Maximisation (EM).

### A. Principal Component Analysis-based spectral analysis

To carry out the analysis, a traffic data profile is created that is based on the features discussed in Section IV, using Principal Component Analysis (PCA). Anomalies are detected based on the difference of the baseline and evaluation profiles. The anomaly detection results are represented by a Receiver Operating Characteristic (ROC) curve, as shown in Fig. 3. Based on a predefined set of anomaly score thresholds, a ROC curve maps a True Positive Rate (TPR) – the rate at which attacks are correctly detected – against a False Positive Rate (FPR) – the rate at which normal traffic is incorrectly identified as anomalous. The goal is to achieve a *high* TPR and a *low* FPR, i.e., be as close as possible to the upper-left of the plot.

We conducted a number of experiments by varying the type and intensity of anomalies, and examining the affect of migration at a source and target data centre. Initially, a smaller dataset was used to determine that the expected behaviour of the anomaly detection approach occurred, with different anomaly types and intensities. This proved to be the case. Subsequently, we ran experiments with the top 200 TCP services, whilst using different virtual service migration intensities, calculated as a percentage of the overall traffic. In these experiments different anomaly types were injected – a DDoS attack, and a vertical and horizontal port scan. Fig. 3 show the results from experiments with a DDoS attack on a source data centre, with an anomaly intensity of 150,000 anomalous flows per five minute period, injected into 600,000 to 800,000 flows. Hence, the anomalies will account for 15% to 20% of the traffic at the moment of injection. Results are shown for anomaly detection performance without (Fig. 3(a)) and with (Fig. 3(b)) migration. Each plot includes results based on a different number of PCA components, ranging from 1 to 8, and denoted by $k$.

As mentioned earlier, Fig. 3(a) shows the results from anomaly detection on an evaluation set of data without migration. Anomalies are detected with 100% TPR for $k = 3$ and higher. For the lower dimensionalities of the analysis model, for a 2% FPR, the TPR is 20% and 40% for $k = 1$ and $k = 2$, respectively. This gives us the baseline performance of the PCA-based anomaly detection technique, with which to compare the impact of migration.

We added the migration component, starting by migrating 2.5% of the traffic. The resulting ROC curve (not shown here for space reasons) looks similar to the one for the baseline traffic (Fig. 3(a)); the TPR values for all $k$ are the same.

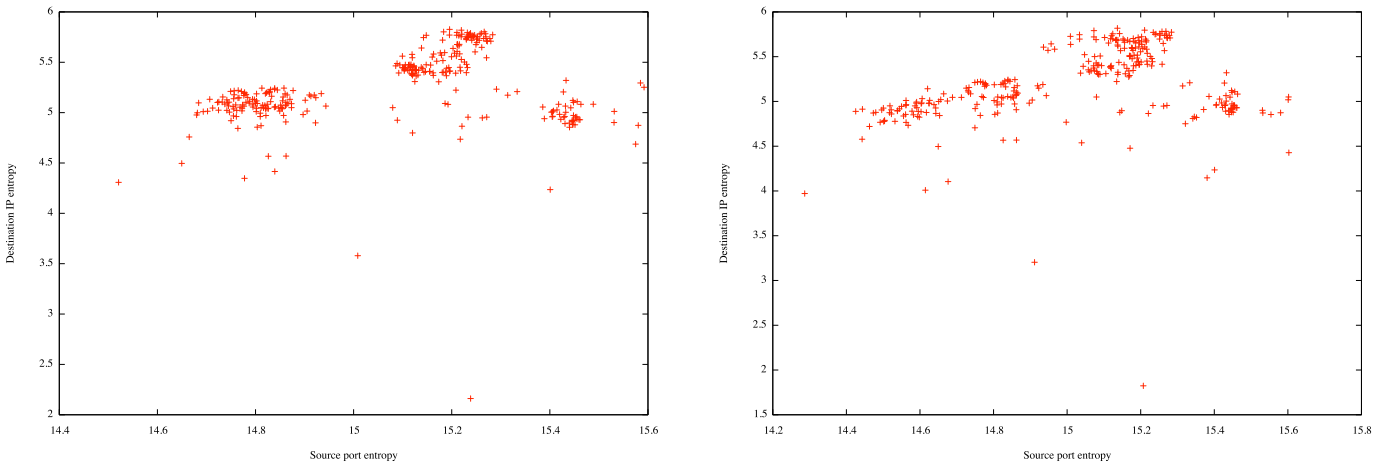


(a) Clustering *before* migration

(b) Clustering *after* migration

Fig. 4. Results of the Expectation-Maximisation (EM) clustering algorithm, plotting source port against destination IP address entropy, before virtual service migration (4(a)), and after (4(b)). A visual shift in clusters can be observed.

For lower dimensions, i.e., for $k < 4$, we observed the anomaly score thresholds differ than for the baseline behaviour. However, we conclude for service migration of this magnitude there is little impact on detection performance.

We conducted further experiments with different magnitudes of migration. For 10% traffic migration, changes in the ROC curves become apparent. For clarity of presentation, we show results from experiments with 20% traffic migration in Fig. 3(b). Migrating this number of services (and its corresponding traffic) could, e.g., be caused by a single failure in a data centre, or represent an aggregate from several smaller migrations, e.g., due to load balancing. As shown in Fig. 3(b), for $k = 2$, improved TPRs can be observed, indicated by the shift in the curve to the upper-left, suggesting an improvement in performance, when compared to the results without migration. However, the TPR is still poor in comparison to higher values of $k$, and improvements are seen at high FPRs. We suggest the detection algorithm should not be used in this configuration. Perhaps the most interesting observation in Fig. 3 relates to the result for $k = 4$, whereby a noticeable drop in TPR can be observed for comparable FPRs, caused my migration, e.g., for a FPR of 2%, a 10% drop in TPR can be seen because of migration; for lower FPRs this phenomenon is more pronounced.

The results show in Fig. 3(b) show good detection performance after migration for values of $k > 4$. In other experiments that use different experimental parameters, e.g., anomaly type and intensity, not shown here for space reasons, we observed instances when this was not the case. Furthermore, we noted this in experiments we conducted using Tsallis entropy, as opposed to the results shown here that are based on Shannon entropy. As expected, we observed similar behaviour at the target data centre. For a further examination of these results, we refer the reader to [15].

To understand the impact of these results on the usability of the algorithm under migration conditions, one must consider the values that underlie the rates in terms of the number of flows that were incorrectly detected. Based on observations from our dataset, the number of flows that are observed for a week could reach approximately one billion, amongst which there are three million injected anomalous flows. For the experiments shown in Fig. 3, with $k = 4$ and FPR of 2%, the successful attack detection rate drops from 100% to 80%. This means that a fifth of the anomalies are missed as being attacks, namely 600,000 anomalous flows. This potentially represents a significant vulnerability to a large cloud data centre and the services it hosts. If we decide to increase the detection threshold, such that the TPR returns to 100%, the FPR increases to 2.5%. This half percent results in false alarms associated with five million non-anomalous flows aggregated into five minute bins – an unusable number for an operator.

### B. Expectation-Maximisation-based Clustering

Another approach to detecting anomalous behaviour is based on clustering, which functions by assigning data points that have similar features to cluster structures. Items that do not belong to a cluster, or are part of a cluster that is labelled as containing attacks, are considered anomalous. There are a number of clustering algorithms available; we chose to base our experiments on the Expectation-Maximisation (EM) algorithm, because it negates the need to pre-define a number of clusters, was shown to give good detection performance in previous research, and its availability in the widely-used WEKA machine learning libraries[3].

In these experiments, we extract the time series entries that have an anomaly score greater than zero, after the spectral analysis discussed in Section V-A. This is done to significantly reduce the number of data points to consider, and to focus on the anomalies caused by attack and migration behaviour. Subsequently, the entries are clustered based on the four entropy-based traffic features, source and destination IP address and port distributions, in order to ignore the fluctuations in number of traffic flows, packets or bytes.

Fig. 4(a) depicts the clusters that are formed without migration; source port and destination IP address entropy features are shown. It appears that three clusters can be distinguished,

[3]http://www.cs.waikato.ac.nz/ml/weka/

including some more dispersed entries. However, the clustering algorithm determines nine clusters, with a log likelihood 1.78. Fig. 4(b) shows what happens after migration. The clusters have dispersed because some of the distributions have been changed after migration. Thus, on this plot half of the data is before migration and half after. It appears as if parts of the clusters have shifted to the right. The clustering algorithm distinguishes only seven clusters, compared to the nine before migration. These clusters are larger and more dispersed, which is confirmed by the log likelihood falling to 0.72. This means that the probability that a selected instance will be put in the correct cluster has decreased more than twice. Even if we apply the algorithm with well-formed clusters after migration, the probability to correctly classify an instance remains low. These results, in a similar to those obtained using Principal Component Analysis, indicate that virtual service migration makes the use of this form of clustering approach unreliable for anomaly detection.

## VI. Conclusions and Discussion

Operators of critical infrastructures are considering moving their high assurance ICT services to the cloud. This implies heightened security requirements; attacks could lead to outages in services that our society depends on, or result in sensitive data being disclosed. An important security measure is to be able to detect when a cloud data centre, and the services that it hosts, are attacked via the network. There are numerous ways to achieve this, including detecting anomalies in network flow summary data – an approach that has seen significant research interest and deployment. One of the essential characteristics of cloud computing is the use of virtualisation technology, which enables services to migrate between different underlying physical infrastructures, both within and across different cloud data centres. Virtual service migration can be used to realise load balancing strategies and improve fault-tolerance to underlying hardware failures, for example. In large cloud data centres, virtual service migration can happen relatively frequently as a consequence of these day-to-day operations.

In this paper, we have examined the affect that wide-area virtual service migration – i.e., migration between cloud data centres – has on contemporary techniques for detecting anomalies in network flow summary data. We have shown that spectral analysis-based detection, using Principle Component Analysis (PCA), and the Expectation-Maximisation (EM) clustering algorithm can be adversely affected by virtual service migration. We argue that, under certain attack and migration conditions, the number of attacks that are missed and false alarms generated by these techniques could render them unreliable and unusable, respectively.

In search of a solution to this problem, we carried out experiments in which virtual service migration behaviour was incorporated into the baseline "normal" behaviour time series. After all, migration of services is arguably representative of the day-to-day operation of a cloud data centre. We found the results of these experiments did not lead to improved detection performance. Further work will examine the reasons why this approach did not yield improved results. Our current thinking about a solution to this problem involves keeping records of the services that are migrated, and using these in a post-detection processing phase to suppress alerts that relate to service migration. We appreciate this approach is not ideal as it requires maintaining and migrating additional state about a virtual service and its clients, which increases overheads and introduces potential privacy issues. Conversely, another approach might involve correlating the alerts from different data centres, in order to determine whether similar behaviour that is indicative of an attack or other problems have been observed.

## Acknowledgement

The research presented in this paper was partially funded by the EC under the FP7 project SECCRIT (Grant No. 312758). The authors are grateful to Dr Bernhard Tellenbach for his support obtaining the PCA-based analysis results.

## References

[1] M. Dekker, "Critical Cloud Computing A CIIP perspective on cloud computing services," white paper, European Network and Information Security Agency (ENISA), Tech. Rep., December 2012.

[2] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Comput. Surv.*, vol. 41, no. 3, pp. 15:1–15:58, Jul. 2009.

[3] B. Grobauer, T. Walloschek, and E. Stocker, "Understanding cloud computing vulnerabilities," *IEEE Security Privacy*, vol. 9, no. 2, pp. 50–57, 2011.

[4] B. Tellenbach *et al.*, "Accurate network anomaly classification with generalized entropy metrics," *Comput. Netw.*, vol. 55, no. 15, pp. 3485–3502, Oct. 2011.

[5] C. Pascoal *et al.*, "Robust feature selection and robust PCA for Internet traffic anomaly detection," in *IEEE INFOCOM 2012*, Orlando, FL, USA, March 2012, pp. 1755–1763.

[6] H. Ringberg *et al.*, "Sensitivity of PCA for traffic anomaly detection," *SIGMETRICS Perform. Eval. Rev.*, vol. 35, no. 1, pp. 109–120, Jun. 2007.

[7] W. Lu and H. Tong, "Detecting Network Anomalies Using CUSUM and EM Clustering," in *Advances in Computation and Intelligence*, Macao, China, June 2009, vol. 5821, pp. 297–308.

[8] I. Syarif, A. Prugel-Bennett, and G. Wills, "Unsupervised Clustering Approach for Network Anomaly Detection," in *Fourth International Conference on Networked Digital Technologies (NDT 2012)*, Dubai, AE, April 2012, pp. 24–26.

[9] D. Brauckhoff *et al.*, "Impact of packet sampling on anomaly detection metrics," in *6th ACM SIGCOMM conference on Internet measurement*, Rio de Janeiro, Brazil, October 2006, pp. 159–164.

[10] P. Mell and T. Grance, "The NIST Definition of Cloud Computing," National Institute of Standards and Technology (NIST), Tech. Rep. Special Publication 800-145, September 2011.

[11] K. V. Vishwanath and N. Nagappan, "Characterizing cloud computing hardware reliability," in *Proceedings of the 1st ACM symposium on Cloud computing*, Indianapolis, Indiana, USA, 2010, pp. 193–204.

[12] M. Nelson, B.-H. Lim, and G. Hutchins, "Fast transparent migration for virtual machines," in *USENIX Annual Technical Conference*. Anaheim, CA, USA: USENIX Association, 2005, pp. 25–25.

[13] T. Benson, A. Akella, and D. A. Maltz, "Network traffic characteristics of data centers in the wild," in *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, Melbourne, Australia, November 2010, pp. 267–280.

[14] D. Brauckhoff, A. Wagner, and M. May, "FLAME: a flow-level anomaly modeling engine," in *Conference on Cyber security experimentation and test (CSET'08)*, San Jose, CA, USA, 2008, pp. 11–16.

[15] K. Adamova, "Anomaly detection with virtual service migration in cloud infrastructures," Master's thesis, D-ITET, ETH Zurich, 2013. [Online]. Available: ftp://ftp.tik.ee.ethz.ch/pub/students/2012-HS/MA-2012-17.pdf