Distributed Architecture

Overview Document

Overview

Large and demanding network infrastructures contain many flow data sources in various locations. Processing large amounts of flow data on a single Kemp Flowmon Collector might be feasible, however this solution does not scale. In a large or expanding network the capacity of a single processing unit will be eventually depleted. Distributed architecture (DA) provides high scalability and load balancing for such demanding environments. Flow data is distributed among multiple units for profile computation and other flow data processing. More units can be simply added to increase both performance and storage capacity. Distributed Architecture provides a central console for management and configuration of all units in remote geographical locations as well as data aggregation and visualization in one place.

Components

There are 3 types of units in DA: Master, Proxy and Slave units. Master and Proxy units are dedicated hardware or virtual appliances. Slave units are traditional Kemp Flowmon Collectors (hardware or virtual appliances). For hardware specification of Master and Proxy units see the Kemp Flowmon Collectors specification document.

Master Unit

Central console for management and configuration of other units. It provides a central interface to all data from all connected Kemp Flowmon Collectors. It provides a web application for data visualization, querying, reporting and analysis. The Master unit gathers data from other units and assembles the final result.
There can be multiple instances of the Master unit. Users work with and perform configuration changes only on the top-priority instance called the Top Priority Master unit (TPM). Slave units, Proxy units and groups are configured on TPM. TPM can initiate data queries on Proxy units and obtain results. Other Master units are synchronized and kept consistent with the TPM and can be set as TPM if the current TPM fails. It is highly recommended to use multiple Master units in a DA topology as a single Master unit cannot be replaced in case of failure. If a single Master unit fails, then with a new Master unit, the whole DA topology must be recreated resulting in complete data loss on all units!

Slave Unit

Slave units are storing and processing assigned part of flow data (see Flow Distribution Models below). More Slave units can be added when needed. Slave units are managed by Proxy units. Slave units can work in two ways as Standalone Slave or Proxy Group Slave. Proxy Group Slave operates as described above (Master – Proxy – Slave deployment mode). Standalone Slave is a Slave unit which operates also as a Proxy unit. In this mode there is only one Slave unit in the Proxy Group (Master – Slave deployment mode).

Proxy Unit

Proxy units are necessary for configurations with multiple Slave units (Standalone Slaves) in one group. A Master unit communicates only with Proxy units to save bandwidth between different locations and for easier firewall configuration. A Proxy unit forwards all its requests to and from Slave units in its Proxy Group. The Proxy unit is used as a single target of flow export (e.g. in one geographical location) and distributes flow data to its Slave units. For groups with a single Slave unit no Proxy unit is needed and Slave acts as Proxy for itself (Proxy Group Slave). One

Proxy unit and one or more Slave units assigned to it creates a Proxy Group. Only one Proxy unit is allowed in a Proxy Group.

Groups

Proxy Group

The Proxy unit and its Slave units form a Proxy Group. Each Slave unit can be assigned to a single Proxy Group only. Proxy Group enables scalability – if the group is overloaded, a new Slave unit can be simply added to take over part of the data and tasks. All Proxy Group Slaves in Proxy Group must be licensed as the same collector model. Only Proxy Groups assigned into the Source Group are able to operate in the DA.

Source Group

A Source Group is formed by one Proxy Group (without High Availability) or more Proxy Groups (for High Availability). All Proxy Groups in a Source Group are identical, deployed in the same location and receive flow data from the same flow sources (hence the name Source Group). If a failure occurs in a Proxy Group and DA has been deployed for High Availability, data collection and query processing will not be interrupted. Proxy Groups are currently not able to recover missing data from other Proxy Groups where the data may be available. As a result, subsequent failures in different Proxy Groups may lead to data unavailability or data loss.

Query Processing

Flow data is stored on Slave units. Master unit stores only aggregated results and metadata. Queries are initiated on the Master unit and forwarded to Proxy units. Each Proxy unit will forward queries to its Slave units. Results from Slave units are aggregated by Proxy if possible and then the results from all Proxy units are sent to Master unit. Master unit then aggregates partial results into the final result provided to the user.

Deployment Modes

Master – Slave Mode

In this mode Master unit communicates directly with Slave units (Standalone Slaves). Each Slave unit is set as a target of flow export for different flow data sources. In the example diagram below, each Slave unit is storing and processing flow data in a different branch office (New York, London and Paris). Master unit provides central reporting and data visualization. Slave units are managed with Master unit from the company HQ.

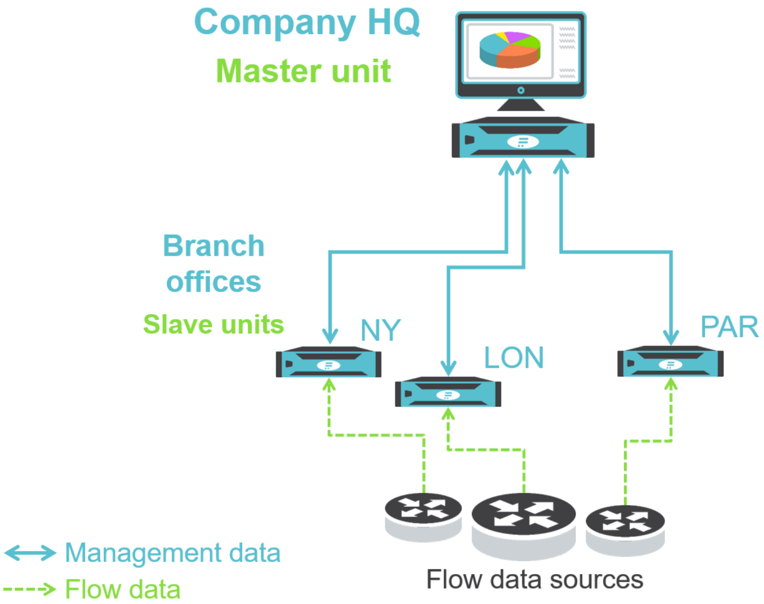

Master – Proxy – Slave Mode

In this mode Master unit communicates only with Proxy units. The Proxy unit is set as a target for flow export and distributes flow data to Slave units (Proxy Group Slaves). Slave units and their Proxy unit form Proxy Group. Slave units can be easily added to Proxy Groups with fully automated provisioning of all the necessary configuration.

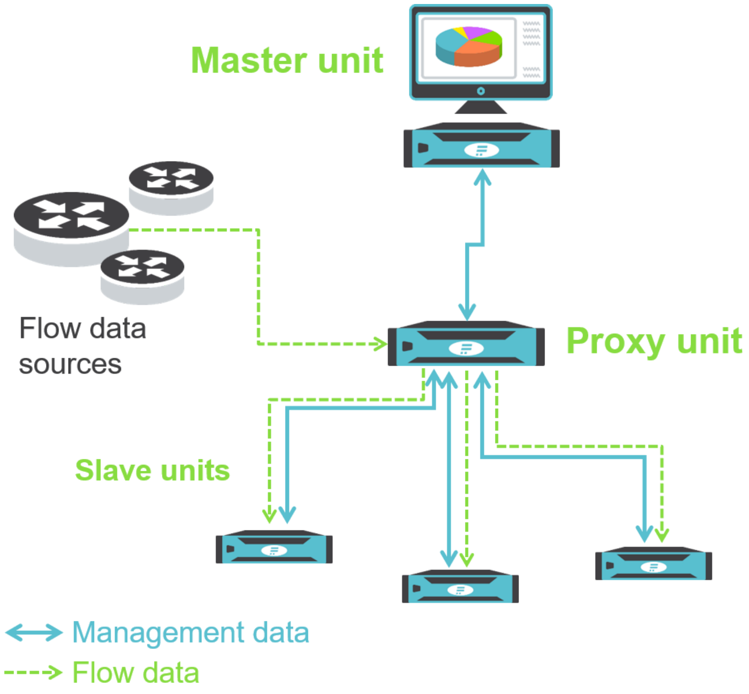

Sample Deployment

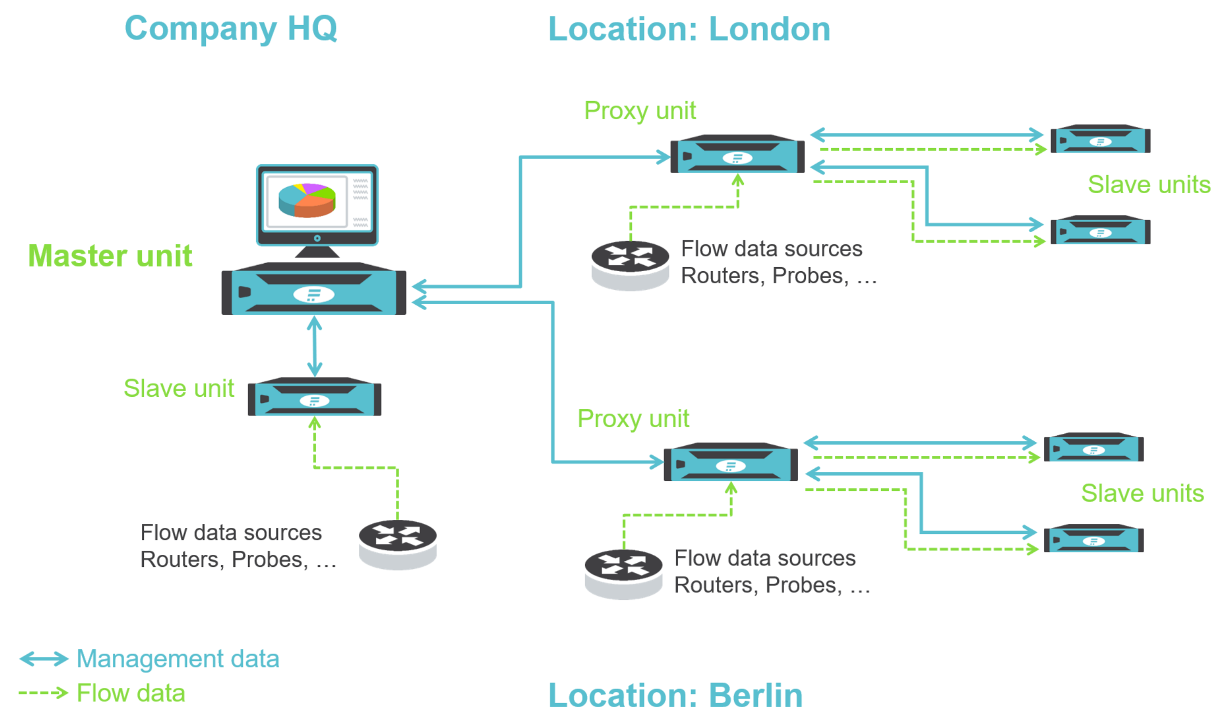

The Company HQ unit is in Master – Slave mode. The Slave unit collects and processes data from flow data sources in the HQ. Remote locations in London and Berlin are in Master – Proxy – Slave mode. The Master unit communicates with Proxy units and those distribute flow data and management changes to Slave units. When the Master unit requests data, it queries the Slave unit in HQ or Proxy units in remote locations.

Flow Distribution Models

This chapter describes flow distribution models, ie. the ways how the flow data can be distributed among Slave units in a Proxy Group. Each approach has certain advantages and disadvantages.

Round-Robin Model

Every Proxy unit distributes flows in a round-robin manner to all Slave units in its group. Incoming flow packets are disassembled, templates are sent to all Slaves units and flows are distributed in a round robin manner. New flow packets assembled by the Proxy unit must have the source IP address of the original flow source.

Advantages:

- Perfect scalability in group ("just add a new device to group")
- All slaves in the group are utilized equally

Disadvantages:

- More complicated data recovery

Flow Source Related

Each Proxy unit maps flow packets from a specific flow source to a specific Slave unit in its group. Incoming flow packets are distributed to Slave units according to flow source address. Flow packets are forwarded as they are.

Advantages:

- Flows from the same flow source are stored on the same Slave unit – they can be used for flow source related detections etc. (e.g. anomaly detection)
- Easy data recovery

Disadvantages:

- Slaves in group are not utilized equally
- Limited scalability - flows from heavy utilized source cannot be distributed to multiple Slave unit

Flow Sources Management

Each new flow source detected on the Proxy unit is reported to Master unit. Master maps this flow source to Source Group, where it was detected and requests primary Proxy Group to obtain flow source metadata via SNMP. Flow source metadata update is requested by Master in a regular manner.

Deleting flow source means to delete a channel of live profile – this is automatically propagated to all sub-profiles and their channels. Delete operation is performed on Master – it will delete a live profile channel (standard profile operation propagated to Proxies and Slaves) and it will also remove flow source from database and from list of flow sources of its Source Group. Deleting the source will discard all its data in the live profile. In sub-profiles, the data will stay intact.

Profiles Management

Profiles are managed by the user on the Master unit. Profile configuration remains the same as in the single system architecture. Selecting parent channels will assign each channel to specific flow sources (as each flow source is representing a root of its channel tree) and hence to Source Groups. The profile is then created/modified on all Slave units in selected Source Groups.
For Flow Sources Related distribution model, profile is created on all Slave units in selected Source Groups as well, even if flow sources are not assigned to all Slaves in a Source Group.

This is necessary in order to:

- keep unified configuration of all Slaves
- allow easy replacement and recovery
- allow changing list of parent channels of existing profile (subprofile of live)

When a profile is created, the affected Source Groups notify all Master units that the profile was created/modified. Every Slave is managing its profiles in the same manner as in non-distributed architecture.

Performance

The performance of a proxy unit is 800.000 fps in deployment with one 10GbE port, and 1.000.000 fps in deployment with two 10GbE ports and flows divided into two listening ports *(note: proxy units with 1GbE ports only have lower performance)*.

Network and Bandwidth Requirements

Communication Port Numbers

DA components use ports 2210 and 4210 for communication between each other. Port number **2210** is used for configuration and control, port **4210** for sending queries and results. Flow data is sent to the listening port of the Proxy unit (set in the GUI), and redistributed to Slave units on port number increased by 30000. For example, if Proxy unit listening port is set to 3000, then the Proxy will forward flow data to Slave units on port 33000). For proper functionality of Distributed Architecture these ports need to be allowed on the firewall.

Bandwidth Utilization of DA Components

Following bandwidth utilization is valid for communication between Master unit and one Source Group. The communication between individual components in Distributed Architecture is not encrypted currently.

Operation	Bandwidth
Normal operations User is working with top statistics, list flow and other data with reasonable result size	**0,1 Mbps**
Large data set delivery User is receiving unlimited list of flows result for large set of data	**10 Mbps**
Major configuration updates exchange	**0,5 Mbps**

Delivery of update packages
Update package transmission between Master and Proxy unit is not considered time-critical and will use only the available bandwidth.

Kemp Flowmon Modules and Distributed Architecture

Kemp Flowmon Modules are software components extending the functionality of the Kemp Flowmon solution with advanced flow data analysis and other features (anomaly detection, application performance monitoring and traffic capture).

Module Deployment

Each module requires a different approach when deployed in Distributed Architecture. This chapter describes these specifics for each extension module.

Following table describes on which units should be each module installed:

	Kemp Flowmon ADS	Kemp Flowmon APM	Kemp Flowmon Packet Investigator
Master unit	✔	✔	✔
Proxy unit	* ✔ / x	x	x
Slave units	✔	x	x
Kemp Flowmon Probes	x	✔	✔

* Kemp Flowmon ADS must be deployed on Proxy units to forward regular Master-Slave communication (events, configuration and queries/responses) when it's not possible to establish direct connection between Master and Slave units. Kemp Flowmon ADS on Proxy units does not do any processing or data storage and does not have any user interface.

Kemp Flowmon ADS

Multiple Kemp Flowmon ADS systems are deployed on multiple Slave units (each Slave needs to be licensed and there has to be the same Kemp Flowmon ADS license on each processing unit). Kemp Flowmon ADS Master unit (GUI) is deployed on Master unit, which provides a central database for all data and detected events from all connected Kemp Flowmon ADS systems. Only one Master unit is supported.

Kemp Flowmon DDoS Defender

Kemp Flowmon DDoS Defender does not support Distributed Architecture.

Kemp Flowmon APM

Multiple Kemp Flowmon APM licenses are deployed on Kemp Flowmon Probes. Kemp Flowmon APM (a regular APM licence) is licensed and installed on the Master unit. Only one Master unit is supported.

Kemp Flowmon Packet Investigator

Multiple Kemp Flowmon Packet Investigator licenses are deployed on Kemp Flowmon Probes and one on the Master unit. Only one Master unit is supported.

Current limitations

- It is not possible to back up flow data to external storage.
- Active Devices are not supported.
- Flow forwarding is supported only on Standalone Slave units.
- Channel options are not supported.
- SNMP live checks of flow sources on Proxy units are not supported.
- AWS Flow Logs processing is not supported.

- A secondary Master unit is not supported when using Kemp Flowmon ADS, Kemp Flowon APM, or Kemp Flowmon Packet Investigator.