Forsway Xtend Hub
Technical Overview

Forsway

Version 4.3, 2020-02-12

Forsway Scandinavia AB, Kanikegränd 3B, 541 34 Skövde Sweden

Table of Contents

1. Introduction

The intention of this document is to provide a detailed introduction to the functionality of Forsway's solution for hybrid Internet services - **Xtend**. The system includes all components that are needed to create a hybrid service. The system consists of the Xtend Hub and the Odin router at the end user's premises. The central components of the Xtend Hub are:

- Mimir NMS
- Mimir Gateway
- Encapsulator/Modulator System (EMS)

The architecture of the system is presented in Figure 1. The following subsections give an overview of the Xtend Hub. The rest of the document includes a detailed presentation of the available features.

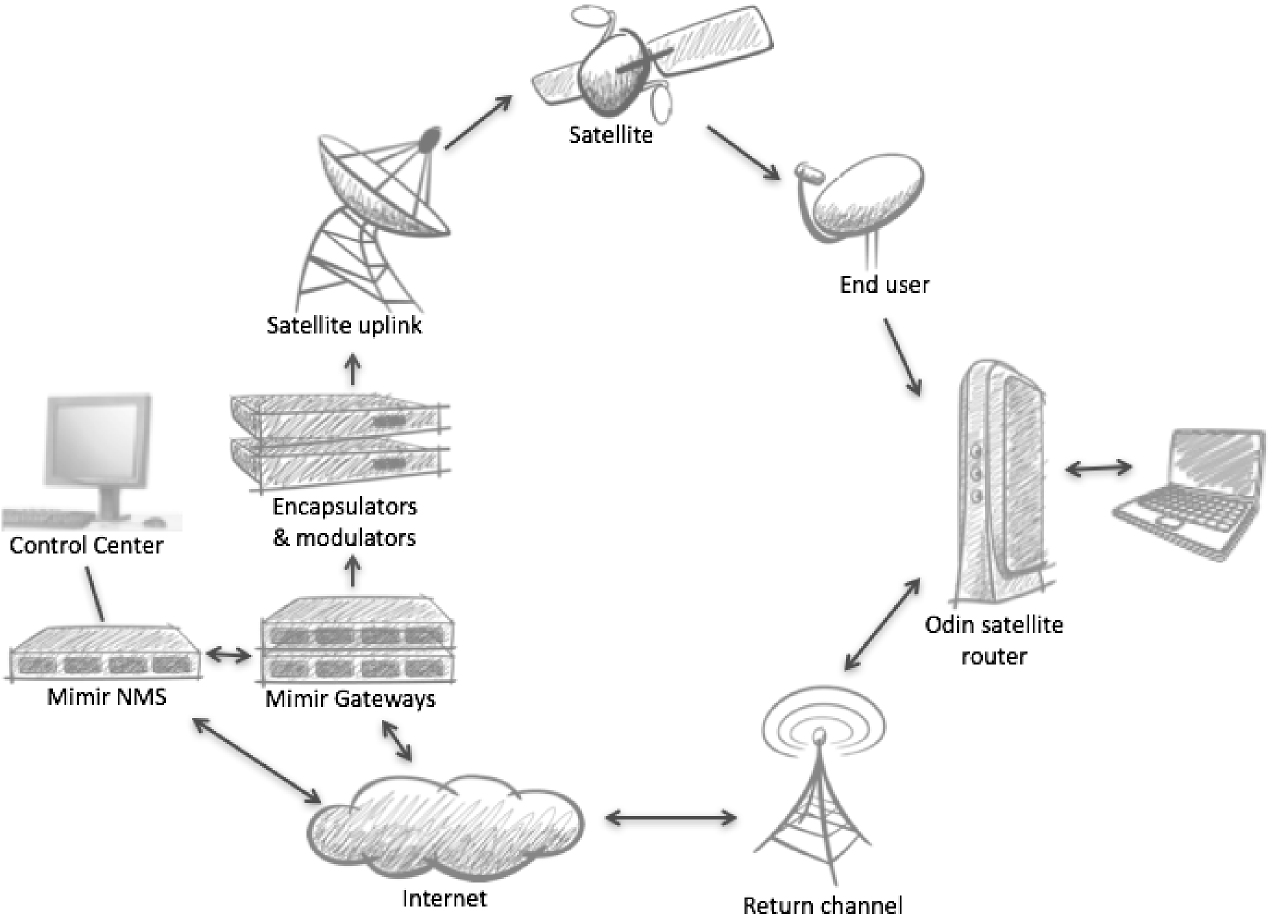

Figure 1. Xtend Hub Overview

1.1. Mimir Gateway

Mimir is a client-server solution for traffic routing and optimizing in a hybrid Internet over satellite service. The server part is basically a proxy server that receives request on one interface, fetches data from the Internet and transmit the fetched data on another interface. The client receives traffic on the satellite interface and transmits on the terrestrial return channel.

Mimir has been designed with two main targets in mind:

- Give a good end-user experience
- Simplify operations and support for the operator

The most important features are summarized below.

Asymmetric Routing

As a hybrid Internet over satellite service uses different transmission paths for the incoming and outgoing data it is important to in a efficient manner support reception of requests on one network interface, while outgoing transmission is done on another network interface.

Protocol Acceleration

The long distance to geostationary satellite has the effect that it takes approximately 240 ms for a signal to travel to the satellite and back to earth. The TCP/IP protocol is not designed for such a long latency and this will affect the possible maximum throughput. Protocol acceleration from a Performance Enhancing Proxy can greatly improve throughput.

On-Demand Satellite Bonding

Mimir can bond medium bandwidth return channels with the satellite service in an efficient way. When bandwidth usage is low all traffic will go over the return channel. When bandwidth demands increases to more than the return channel can provide, satellite is automatically activated and bonded with the return link to give the end-user the combined bandwidth of the satellite link and return channel. Latency sensitive protocols are always prioritized for the low-latency return channel.

Link Aggregation

The system supports different type of link aggregation, for example by combining the bandwidth of two mobile connections for improved uplink performance.

Terminal Auto Configuration

All satellite parameters are automatically sent to the terminal when connecting. The end user only needs to configure the return channel.

Security

The Mimir server allows secure authentication at login time from the user of PKI certificates. The certificates are preinstalled in the Odin satellite modems from factory and removes the burden of username and password management. All communication between the Mimir server and terminal can be encrypted.

Multi-Transponder

Mimir supports multi transponder environments with automatic load balancing between transponders, based on geographic properties of end users.

Virtual Network Operators

Several virtual network operators can be defined in the system. The virtual network operators manage their own users and services plans.

User Management

Mimir NMS includes a system for simple user provisioning, allowing the system administrator to configure and manage users, terminals, and different types of accounts and services. Third party provisioning and billing systems can be integrated through a provisioning API.

Simple Installation

The Mimir server is preferably deployed at the uplink site. The Mimir server is most commonly preinstalled to simplify deployment. The Odin terminal can be factory configured for a specific operator.

To make deployment easier the Xtend Hub is in many cases bundled with an integrated encapsulator/modulator.

1.2. Mimir NMS

Mimir NMS manages the whole system. Mimir NMS simplifies terminal management in a multi transponder configuration. Terminal management and configuration can otherwise quickly become an overwhelming task in large system with multiple transponders, potentially covering different satellite footprints.

Mimir NMS is the central connection point of the system. All terminals first connect to the Mimir NMS, where they are dispatched to a Mimir service based on population and service congestion. Once connected to the service all tuning parameters required to access the satellite service are pushed out to the terminal. The populations used for load balancing can be based on geography, i.e. which satellite footprint, or used to define virtual operators.

Mimir NMS keeps track of the health of all components in the system and can reconfigure the system to use redundant devices when a failure occurs.

Mimir NMS can be deployed either on servers installed locally in the operators LAN or provided as a cloud based service operated by Forsway.

1.3. Encapsulator/Modulator System (EMS)

The Xtend Hub have integrated support for some specific encapsulators and modulators to allow the Mimir server to control and monitor some important parameters.

The devices that are currently best integrated are:

- Forsway EMS
- Tebkom ODG-200
- TeamCast SmartGate

- Ayecka ST1
- Encapsulators and IP modems from Work Microwave

The most important integration points are listed below.

IP-MAC Table Management

Mimir can manage the IP-MAC association table in the encapsulators. This means that dynamic associations can be used allowing the system administrator to configure IP pools and let the system manage IP address assignments.

The supported encapsulators can either manage thousands of entries in the IP-MAC association table or support on-the-fly IP-MAC association.

ACM

The supported encapsulators and modulators can be equipped with ACM controllers. The OpenTMP protocol is used for reporting signal quality. One important feature of this protocol is that the terminals can be intelligent and request MODCOD changes when needed. Compared to many other ACM protocols this means that signal quality messages can be sent with a lower frequency.

Forsway's Odin satellite modem supports sending ACM messages in the OpenTMP protocol. Mimir can tunnel the messages to an ACM controller located on the M&C network, to avoid making the ACM controller available to the public Internet.

Mimir can also read out signal quality from the ACM messages and presents the signal quality for all online users in the web based user interface.

Bandwidth Monitoring

Mimir can read out the currently available bandwidth from the encapsulators. This is important for some types of ACM configurations, where the bandwidth constantly changes.

VRRP

For encapsulators supporting redundancy with VRRP the Mimir Gateway can automatically switch to a redundant encapsulator in case of a failure. This makes it possible to achieve a very high availability.

2. Core Features

2.1. Performance Enhancing Proxy

Due to the long distance to geostationary satellites, the time for a signal to make a jump over a satellite is 240 ms. The TCP/IP protocol is not designed for such a long latency and this will affect the possible maximum throughput. Mimir solves this by including a performance enhancing proxy (PEP). The PEP splits the communication link, first in the Mimir Gateway and secondly in the receiving terminal and uses protocols optimized for the jump over satellite.

The TCP protocol options (congestion control algorithm, TCP window size etc.) used from the server through the satellite uplink are optimized for the latency typical in satellite communication. This way the throughput can be substantially increased compared to using the default TCP options.

Figure 2 illustrates the principles of the PEP. The connection between the end user's computer and Internet server is split in three parts. When the Mimir server receives a request from the end user it will send the request to the correct Internet server, but unlike the normal mode of operation, when a response is received from the Internet server it will not wait for an acknowledgement from the end user's computer, but will act as an end-point on the terrestrial connections to make the remote server believe that the packets have arrived at the destination. In next step a new connection is opened over the satellite, using a protocol that is optimized for satellite communication. In the client terminal the connection will once again be terminated, and a new connection is opened against the computer using standard protocols.

When using VPN based on protocols such as UDP or IPsec, TCP traffic in flowing over the VPN won't be accelerated.

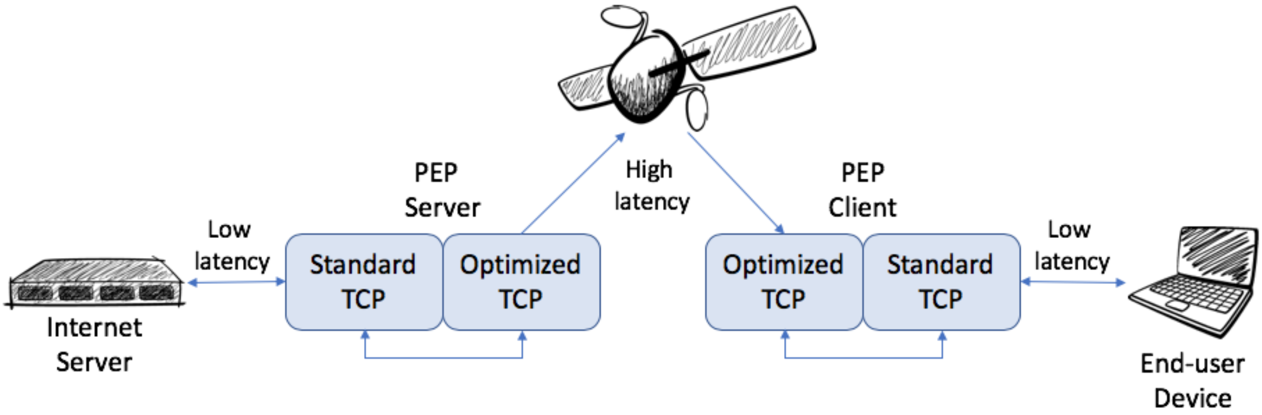

Figure 2. PEP Principle

2.2. Asymmetric Routing and Tunneling

One of the main principles of Mimir is the support of asymmetric routing where requests from a terminal can come on one interface and responses sent on another. To achieve this in a transparent way to applications the Mimir service acts as a proxy through which all traffic is routed.

For the proxy to be usable to the client terminals, the terminals need to route all traffic through the underlying network infrastructure to Mimir. Remember that the terminals will be using different terrestrial connections, often provided by third party/independent operators.

The method employed is an IP-tunnel/VPN solution based on the 'tun – tap' driver interface available in Linux. This allows the system to set up a network interface and default routing that make the normal connections always routed through the 'tun' interface.

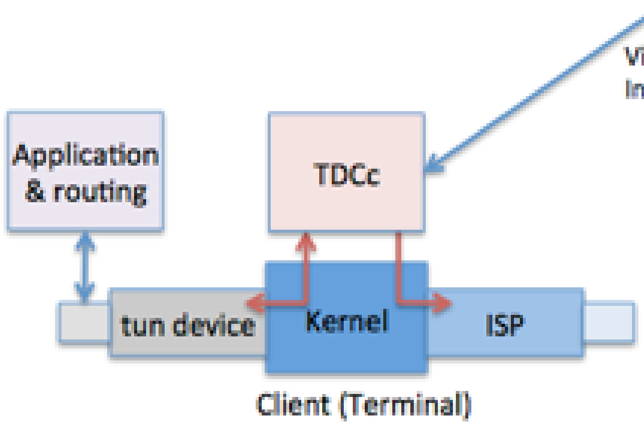

Figure 3. Tunnel Principles

When the 'tun' interface is configured it should have an IP-address on a different network than the IP-address of the terrestrial return channel, not to collide with IP numbers given to the terminal by the network operators on connection. The IP-address must also be on a different network address than the LAN interface. When the connection is initiated the terminal needs to be registered and authorized to connect to the proxy server.

The client terminal runs a tunnel/VPN client program (Tunnel Control Daemon –TCDc, see Figure 3) that at connection time assist in the authentication of the client terminal based on the VPN log-in user settings or pre-installed certificates/keys in the client terminal.

The tunnel/VPN solution includes a control channel for each client terminal that only sends traffic on the terrestrial return channel. The purpose of the control channel is to keep the connection alive, to detect if a client terminal is no longer connected or to inform the client if bandwidth limitations have been reached.

The connection to the return channel ISP will normally use the ppp-protocol to set up the connection over the terrestrial modem, be it PSTN, GPRS, 3G, or other.

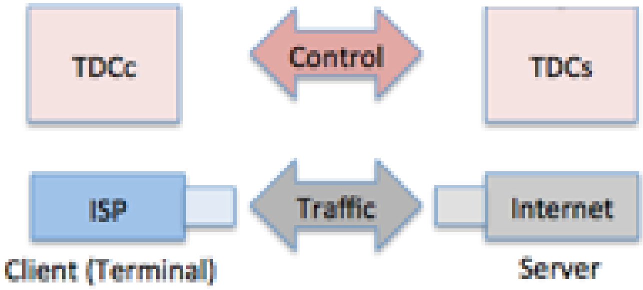

Figure 4. Tunnel Communication

2.3. On-Demand Satellite Bonding

Mimir supports on-demand downlink bonding of satellite and return channel bandwidth. This means that the satellite link and return channel can be combined in an efficient way to improve user experience and decrease the usage of satellite bandwidth.

On-demand bonding is especially useful for medium bandwidth return channels with low latency, e.g. slow ADSL connections.

When bandwidth usage is low all traffic will go over the return channel. When bandwidth demands increases to more than the return channel can provide, satellite is automatically activated and bonded with the return link to give the end-user the combined bandwidth of the satellite link and return channel. When bandwidth demands decrease again the bonding is automatically disabled.

Bonding with a satellite link has the impact that Internet traffic is affected by satellite latency. To give the best possible user experience only bandwidth heavy services that are not very latency sensitive are bonded: web surfing, file downloads, and video streaming. All other services are treated as low-latency services and are routed over the low latency return channel.

2.4. Link Aggregation

Mimir supports advanced link aggregation, allowing the terminal to combine the bandwidth of several data links. The functionality can be configured in several ways. Example scenarios:

- Combine the downstream bandwidth of a satellite forward link with several combined return channel links, and combine upstream bandwidth of the return channel links.
- Use only the satellite forward link downstream, but combine several return channels for increased upstream bandwidth.

Should a link be lost it will be automatically detected and removed from the aggregation group. If it reconnects it is added again.

2.5. Multi-Transponder

One physical server can provide up to four separate Mimir services. A service is the complete set of functions to provide users Internet access on a specific PID. A service is always using a specific IP encapsulator. What physical Mimir Gateway the service runs in is less important, at startup time the service will be assigned to run on one of the Mimir Gateways connected to the correct encapsulator.

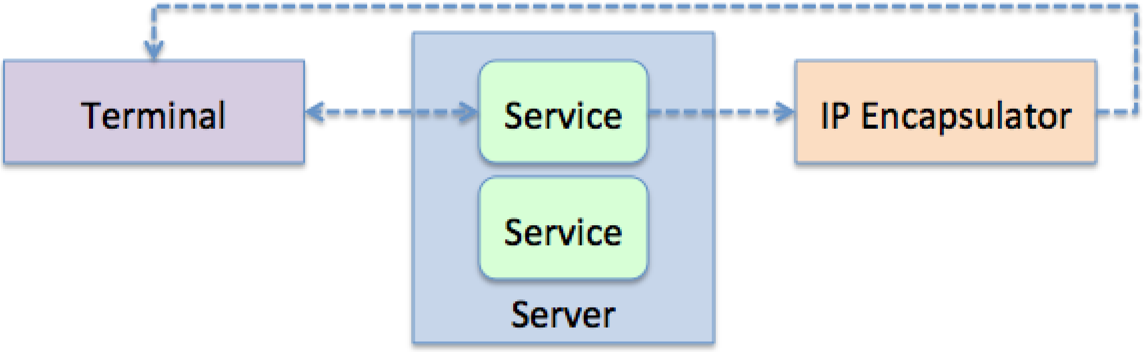

Figure 5. Concept of server/service

Populations are a way to group users based on some property, most commonly geographic location. Other properties can also be used, i.e. to differentiate between virtual network operators. Services are assigned to populations. Several services can share one population.

When a user connects to Mimir NMS it will identify to which population the user belongs and dispatch the user to the least congested service for the user's population. Once connected the satellite parameters for the service will be sent to the user's terminal for automatic configuration.

2.6. Service Plans and QoS

Service plans allow the operator to set limits on a user's maximum bandwidth (MIR) and number of gigabytes that can be downloaded per month. Mimir will keep track of the connections for each client terminal and monitor usage and apply the limits. Fair access policies (FAP) can be defined to let users' available bandwidth be limited in several steps when monthly data usage thresholds are reached. The FAP can also include limitations based on the congestion of the service.

The CIR of a user is dynamically calculated by comparing the MIR of a user's service plan in relation to all other logged in users, after applying Data and Policy FAPs.

The FAP is a means of dividing the bandwidth fairly among the users by limiting the available bandwidth for heavy users. The FAP is implemented with limitations on three levels based on different input:

Level 1 - Data
Example:

- Start with 100% of service plan bandwidth (MIR)
- When 50% of quota is used, limit MIR to 75%
- When 75% of quota is used, limit MIR to 50%
- When 90% of quota is used, limit MIR to 10%

Level 2 – Congestion
The congestion FAP limits can limit bandwidth for groups of users when the service is crowded. This is a way to let for example business users be less affected by congestion then end-consumers.

The target maximum number of users can be set for a service. This value is used for calculating congestion levels. The value is only used for calculations and does not block users when the user count is reached.

Based on the max user value a number of congestion levels are extracted. The value in parenthesis indicates when stepping down from the level above.

1. Not congested < 60% (55%)
2. Congested 60-75% (70%)
3. Very congested 75-85% (80%)
4. Extremely congested > 90%

The bandwidth limitation can be specified for each congestion level, e.g.

- Not congested, give 150% bandwidth
- Congested, give 100% bandwidth
- Very congested, limit to 75% bandwidth
- Extremely congested, limit to 50% bandwidth

Level 3 – Traffic Type

The Protocol FAP is based on the complete bandwidth usage of the whole service, limiting the total bandwidth to fit the available satellite bandwidth.

Different types of traffic is throttled in different ways to make sure the user experience continues to be as good as possible when bandwidth is throttled. Network traffic is divided into three types.

1. TCP-bulk transfer (e.g. downloading files from Internet).
2. UDP traffic and small TCP traffic (VoIP, SYN handshake, ACKs etc.)
3. All the rest.

Traffic type 2) has highest priority and there is always 10% of the bandwidth reserved for this, i.e., traffic like VoIP and TCP handshake will always have higher priority and will always have at least 10% of the service bandwidth. It can consume more bandwidth if needed. Traffic type a) has lowest priority, i.e. bulk transfers will always be throttled first if the bandwidth is strained.

If the total bandwidth goes over a specific value the traffic type bulk transfer will be limited to a first limit (e.g. 90%). If the bandwidth usage is still too high, the bulk transfers will once again be reduced (to 80%). When bandwidth usage drops bulk transfers will once again be given more bandwidth.

Applying the FAP

The FAP penalty for each user is calculated as follows:

User bandwidth = Subscribed MIR * Data penalty * Congestion penalty

The Protocol FAP is in the next step, where the bandwidth is limited in an equal way for all users.

FAP limitations are recalculated every 2-3 seconds.

2.7. Virtual Network Operators

The system support multitenancy with the possibility to define several Virtual Network Operators (VNO). Each VNO have full control of their users and service plans. All features of the system such as multi-transponder load balancing and ACM can be used with VNO.

The Mimir API allows for different VNOs to use different billing systems.

2.8. ACM/VCM

Mimir supports the OpenTMP ACM messaging protocol. Terminals send the ACM messages to Mimir. Mimir forwards the message to the ACM controller. This way the ACM controller does not need to be available on the public Internet. In the forwarding process Mimir reads out signal quality information from the ACM messages and displays signal quality for all connected terminals, a perfect help for finding misaligned antennas and user support. The possibility to display signal quality is available even when ACM is not used for the transmission, as long as the terminals support the OpenTMP protocol and DVB-S2 is used.

The possibility to assign users to different geographic populations makes it easy to support VCM configurations, where different users are located in areas with different reception quality. A separate population can be created for each MODCOD.

2.9. Header Compression

The system supports header compression on the return channel, saving about 20 bytes per packet. This can reduce the amount of traffic on the return channel with 5%-20% depending on usage pattern Typical savings for general office use is about 10%.

2.10. Robustness

The system supports high availability through redundancy and clustering. Mimir Gateway servers, Mimir NMS servers, and encapsulators can be redundant.

Mimir Gateway

A cluster of Mimir Gateways can be managed by the Mimir NMS. Any Mimir Gateway that is connected to the same encapsulator as a failed gateway can automatically take over service, with the restriction of maximum four services per Gateway. It is possible to have hot standby backup Gateways that will only be used in case of a failure. Once redundant Mimir Gateway can act as backup for four services, i.e. up to four primary Gateways. When the NMS detects a failed Gateway the services on that Gateway are reassigned to a backup Gateway within in a few seconds.

Mimir NMS

Mimir NMS can be provided either as a cloud based managed service operated by Forsway or on locally installed at the service operator’s facilities, operated by the service operator.

The cloud based managed service is based on the high availability of Amazon AWS cloud services. A hot standby NMS server is always ready to take over in case of a failure. After a failure a new server will automatically replace the failed server to regain redundancy.

For local deployments the Mimir NMS can be installed in a high-availability cluster configuration. Terminals connect to a virtual IP address, shared between the servers in the cluster. Should the active Mimir NMS server fail, another server in the cluster will quickly takes the active role. A NMS failover typically require no more than a few seconds.

Users already connected to a Mimir gateway will not be affected by the NMS failover. The Mimir Gateways will continue to operate even though they lose connectivity to the NMS for a short while.

Encapsulator/Modulator Systems (EMS)

Encapsulator redundancy is supported through VRRP. The Mimir server monitors signals to switchover data output to a redundant device in case of a failure. The encapsulators need to be configured for redundancy. The redundant encapsulator has to monitor the primary in order to in case of a failure switch over to active and request the RF switch to change input.

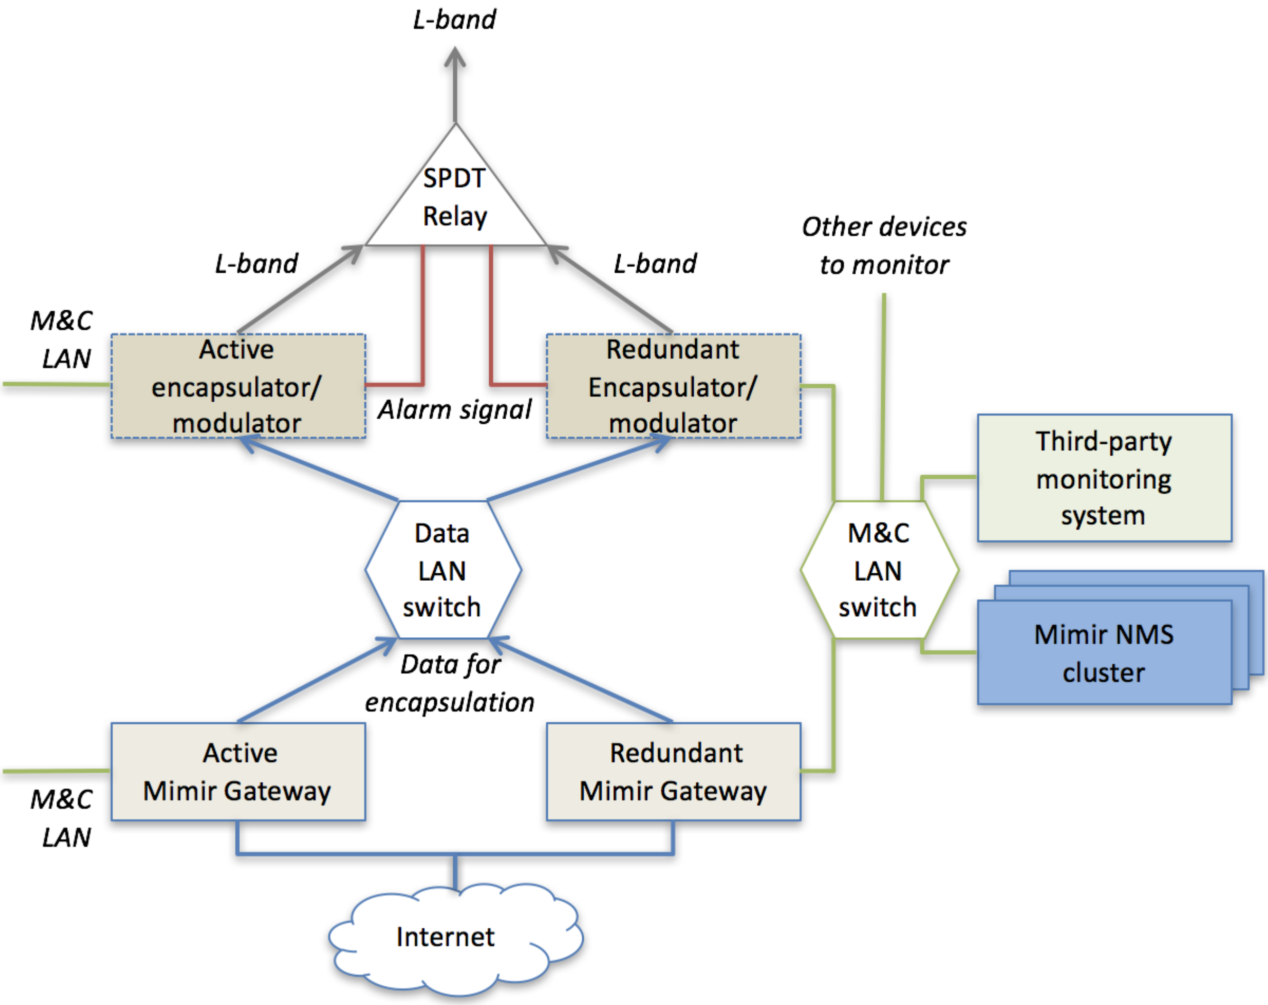

Figure 6. Redundancy Overview

2.11. Public IP Address

The default behavior for Mimir is to hide all terminals behind a NAT. It is possible to assign public IP addresses/subnets to be assigned to specific terminals. The endpoint for a public IP address can be either the terminal or a computer connected to the LAN port of the terminal.

2.12. Web Cache

The Mimir server includes a web cache (Squid). The web cache can reduce bandwidth usage on the Internet connection to the satellite uplink by caching popular web pages. The caching can also help in improving user experience through reduced latency, as commonly accessed web pages are already available locally in the server.

2.13. Prefetching

The Mimir server includes support for prefetching. Prefetching means that when a user loads a page, linked pages will be automatically downloaded in the background. This can help reducing latency, but at the expense of increased bandwidth. The default prefetching configuration will only load text and html content to minimize bandwidth usage. Note that prefetching is only made in the server at the satellite uplink, not in the end user’s terminal.

3. Security Features

Mimir is equipped with a number of security features. Some are used for the security of the individual user and some are required by legal authorities.

This section describes the inbuilt security features of Mimir and how Mimir can be used together with external systems to provide additional security features.

3.1. Secure Login

The secure login procedure for Mimir is based on public and private certificates, commonly referred to as Public Key Infrastructure (PKI). PKI uses three kinds of certificates to create a platform for secure authentication

- Master certificate (CA)
- Private certificates unique for each Mimir Gateway or Odin router
- Public certificates unique for each Mimir Gateway or Odin router

Forsway keeps the master certificate (CA). This is used for generating the other certificates. Private and public certificates for a Mimir server are generated before installing the server. Private and public certificates for each Odin terminal are generated as part of the manufacturing process and stored on the device. The client certificates are based on the MAC address of the terminal.

Mimir login procedure:

1. When the terminal has connected it will exchange public certificates with the NMS or Gateway
    a. Mimir verifies the terminal's public certificate with the Mimir private
    b. The terminal verifies Mimir's public certificate with the terminal's private
2. The terminal's public certificate contains information about the MAC address. Mimir uses this to check if there is a user assigned to this MAC address
3. If there is a user assigned to the MAC address the terminal is allowed to log in

As the client certificates are generated during manufacturing the end user or administrator won't need to state login/password for each terminal.

To assign a terminal to a user the administrator only needs to enter the MAC address of the terminal in the Mimir NMS or in a integrated third party provisioning system.

Certificate Properties

The certificate mechanism for Mimir is using Secure Socket Layer (SSL) for certificate generation and verification. The OpenSSL library is used to provide SSL functionality. OpenSSL is a robust, commercial-grade library that provides cryptographic functionality. OpenSSL is used on many Internet servers and several of the top websites. It is validated under the FIPS 140-2 computer security standard by the National Institute of Standards and Technology's (NIST) Cryptographic Module Validation Program (CMVP)

The certificates have the following properties:

```
Public Key Algorithm: rsaEncryption
RSA Public Key: (2048 bit)
Signature Algorithm: sha1WithRSAEncryption
```

3.2. Encryption

Mimir includes support for encryption of all communication channels between the end user’s terminal and the Mimir server. Encryption is configured for each service plan. Different plans with different encryption can run side by side. One plan can have encryption disabled, a second use a high-speed encryption, and a third a very secure encryption.

There are three different connections between the terminal and Mimir, of which all can be encrypted

1. Control channel. The control channel is encrypted with standard SSL encryption using the AES algorithm. The login certificates described in section 3.1 are used for securing the SSL connection.
2. Satellite forward link. The forward link is encrypted with the algorithm selected for the user’s plan.
3. Terrestrial return channel. The return channel is encrypted with the algorithm selected for the user’s plan.

Encryption Algorithms

The two algorithms AES and SOBER are supported. The AES cipher can use 128 bit and 256 bit keys. The SOBER cipher can use 128 bit keys. AES is a type of symmetric cipher, which has high security level, but a relative slow speed. SOBER is stream cipher with high encrypting speed.

When terminals connect the SSL encrypted control channel connection is enabled. The encryption for the data link uses the hashed SSL keys as main encryption key, which means further key exchanges are not necessary.

Performance

The performance of encrypted transmissions is mainly limited by the decryption capacity of the Odin satellite router. In most scenarios the performance of the most secure cipher, AES256, is good enough.

- SOBER128 algorithm gives a throughput of 30+ Mbit/s on Odin (F-30)
- AES256 provides the best security and gives a throughput of 10+ Mbit/s on Odin (F-30)

The use of encryption has a small bandwidth penalty of 8 bytes per IP packet.

Implementation Details

The Mimir server runs encryption in user space. The terminal has a specific kernel module to achieve best possible throughput.

CTR mode for AES cipher is used to encrypt/decrypt a message.

All ciphers (AES-CTR and SOBER) need an initialization vector (IV) for every message. The IV is also a nonce (number used once), which is used to ensure that the key-stream is different for different messages encrypted with the same key. In our implementation, the IV is a uint64_t message number which increases from 0 to the 2^64 on the Mimir server. On the client side it will decrease from 2^64 to 0. Currently, the IV will be inserted to the beginning of the payload after encryption and removed after decryption (The IV should always be plain text). Neither server nor client will check whether the IV is valid. The lack of IV/msg number verification will potentially increase the repeat attack risk, but it is possible to transmit 2^32 messages with valid and non-repeated IV for the forward link, which is considered to be secure enough.

3.3. Content Blocking

Authorities sometimes require ISPs to block specific content, most commonly certain websites. Many countries requiring content blocking have a national firewall through which all traffic need to pass, creating problems when uplinking from another country. Some countries require each ISP to block certain content.

Content blocking is not part of Mimir but instead provided by other systems. Recommended configurations for the most common scenarios are described below:

1. National firewall and in-country uplink. All traffic on the national Internet backbone is affected by the national firewall and nothing extra needs to be done.
2. National firewall, but uplink from another country. This can be solved in several ways. Note that all solutions require dedicated high bandwidth from the satellite uplink into the country where the firewall is located.
    a. Let the Mimir server use a national proxy
    b. Use a VPN from the Mimir server to a national gateway
    c. Use a leased line into a national gateway for the Internet connection to Mimir
3. ISP responsible for blocking. In this scenario, a firewall with content filters need to be placed on Mimir’s connection to the Internet. Third party firewall software can also be installed on the Mimir server.

3.4. Connection Logging

Many counties require the ISP to store logs that can be used to trace which user that is assigned to a specific IP address at a specific time. One example of such regulation is the Data Retention Directive in EU.

The situation gets more complex by the fact that Mimir uses NAT to let all logged on users share the IP address of a Mimir server. The effect of this is that for the outside world all users appear to come from the same IP address. This is solved by the connection logging functionality in Mimir that logs every access to the Internet for each terminal. The log can later be used for traceback of a connection to a specific customer on request from legal authorities.

3.5. Lawful Interception

Lawful interception (LI) is the process when legal authorities collect and analyze traffic data in real time to prevent crimes and terrorism. There are European and American standards for lawful interception and the specific implementation will be slightly different depending on the used standard.

Lawful interception wiretapping is done on demand and should be transparent to the law enforcement agency (LEA) and the operator. There are three different implementation techniques of LI:

- Active: Direct local interception, either support for LI in the network device or support for it in a connected external device.
- Semi-active (hybrid): Traffic is sniffed passively but with interaction with the Mimir server to be able to perform IP-address/customer matches. Active interception is then made for the specific target customer.
- Passive: No interaction with ISP required, only interception point for LEA device. Traffic can be sniffed on the ISP Internet connection and then analyzed passively offline.

Lawful interception is is done by systems external to Mimir Gateways. To add this support all data to and from Mimir will have to pass through a switch with port mirroring to a LEA device. To understand which data is for which user the LEA device will need to interact with Mimir.

4. Mimir NMS

Mimir NMS includes a web portal for service and user provisioning, configuration, and monitoring. An overview of the portal is presented here. Detailed usage instructions can be found in the operator manual.

4.1. Service Configuration

Service configuration is a three-step process.

1. All encapsulators in the system are added, including the tuning parameters required for auto configuration of terminals
2. A new Mimir Gateway is added. The Gateway itself will first have to be configured to connect to Mimir NMS, next it is added to the system through the Mimir NMS. All encapsulators to which the server has physical access are added as part of the Mimir Gateway configuration.
3. A new service is added. The service definition includes what encapsulator should be used, on what PID it is running, and which population it belongs to.

When everything is ready the service is automatically started on one of the Mimir Gateways connected to the selected encapsulator.

4.2. Service Plans

Service plans consist of a fair access policy (FAP) and a plan.

The plan defines the user’s type of service, maximum bandwidth, monthly quota, and encryption algorithm. Service type can be hybrid, on-demand bonding, or link aggregation.

The FAP defines at what percentage of quota and congestion bandwidth shall be reduced, or increased, by a specific percentage. The FAP and plans are described in detail in section Service Plans and QoS.

4.3. User Provisioning

Mimir NMS provides an easy to use user provisioning interface. Users can be added, removed, disabled/enabled.

To be able to connect a user needs an assigned terminal. The only thing needed to bind a terminal to a user is to enter the MAC address of the terminal. When later connecting with the terminal it will automatically login as the specific user, with the correct service plan. See section 3.1 for details about the secure login process.

Third party billing and provisioning systems can be integrated through the Mimir provisioning API.

4.4. Monitoring

Mimir NMS can visualize the system configuration, with current usage and health status of all components. This is an easy way to keep watch over the whole system.

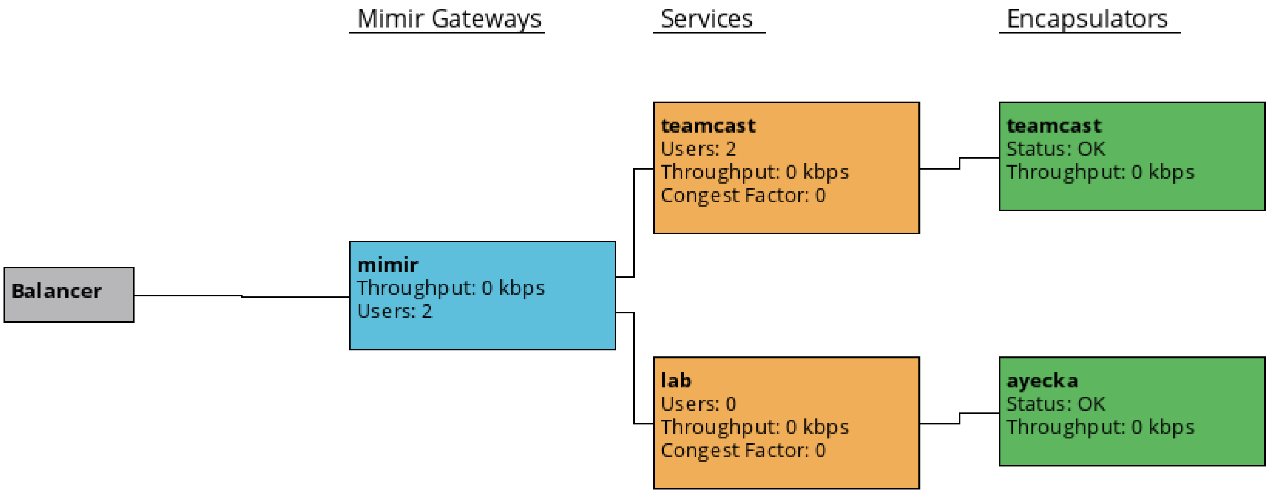

Figure 7. Visual presentation of system configuration and health

It is possible to visualize system usage and behavior in different ways:

- Online users, with information on satellite signal quality, current bandwidth, geographical location, and the possibility to remotely control users' terminals
- For services it is possible to show graphs with history of upstream and downstream bandwidth and login count
- For encapsulators it is possible to show graphs with bandwidth history
- For users it is possible to show connected sessions and graphs with bandwidth history
- Application and usage is tracked with Deep Packet Inspection (DPI) and can be visualized for users, services, and encapsulators
- Detailed system logs can be displayed as a help to trace issues

4.5. VNO Management

Virtual Network Operators (VNO) are created in the Mimir NMS. VNOs are assigned to populations to control to what services connecting terminals are assigned.

4.6. Administrator Management

Mimir NMS includes support for several types of administrators with different login/password.

- System administrators have full control of all parts of the system, including the creations of VNO and VNO administrators.
- VNO administrators have access only to one VNO. The authorities of VNO administrators can be further customized on individual administrator basis so that some might have full access to

modify and some might only have read access to some or all parts of the system.

4.7. Terminal Remote Control

The remote terminal control feature in Mimir NMS is a good help for end user support. The support engineer can open the command line of any connected terminal for error tracing. The only requirement for remote control to work is that the return channel and connection from the terminal to Mimir is working, making remote control possible even if there are problems with the satellite link.

Remote control of an Odin router is initiated by pressing the connect button in the list of online terminals. This will open a new web browser window with command line access to the terminal.

4.8. Terminal Geolocation

All terminals report information on mobile base stations and visible Wifi networks to Mimir. This information can be used for calculating an approximate geographical location of the terminal. The approximate location can be displayed on a map for helping support engineers identify the cause of satellite reception problems.

5. Software Components

The Mimir hub consists of a number of software building blocks:

- Mimir Gateway, manage user sessions and IP traffic routing.
- Mimir NMS, NMS responsible for managing the cluster of Mimir Gateways and assigning users to Mimir services. Controlled through a web based portal.
- Mimir Management API, provides an interface for user provisioning through third party systems.
- Database backend, stores user credential, traffic accounting data, IP pools etc. Mimir uses the MySQL database.
- Squid web cache, used as web cache and for prefetching.

Figure 8 depicts the components that the Mimir system consists of. The components can initially be on the same hardware, but when the system grows and the capacity needs to be increased the main systems, i.e. the Mimir Gateway and Mimir NMS should preferably be run on dedicated physical servers.

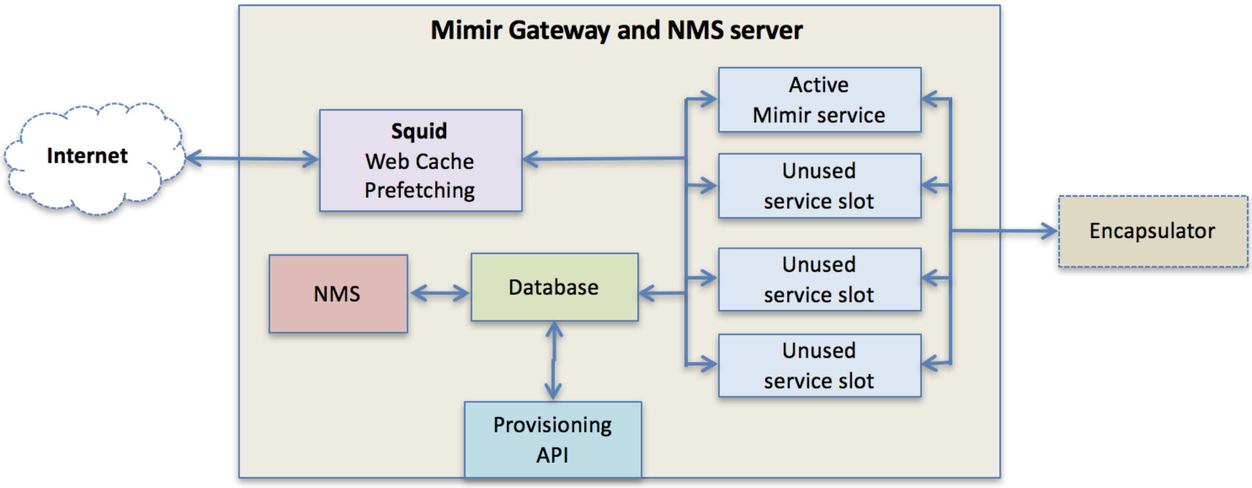

Figure 8. Building blocks in a minimal deployment

6. Installation of Mimir

Preinstalled servers with Mimir NMS and Mimir Gateway are in most cases easy to deploy in a teleport. A short overview is presented here, with detailed instructions in the operator manual. Some preconditions must be fulfilled:

- Direct connection to the Internet, preferably with a fixed public IP address for each server
- IP encapsulators and modulators available

Figure 9 shows an overview of the functionality divided over Mimir and other devices at the satellite uplink. The box ”Mimir NMS and Gateway server” shows all Mimir software running on a single server, e.g. for a demo. Most commercial installations use separate servers.

All Mimir Gateways in an installation share the same Mimir NMS server.

It is possible to run up to 4 Mimir services on one server. Each service can be run on a separate encapsulator or as a specific PID on the same encapsulator.

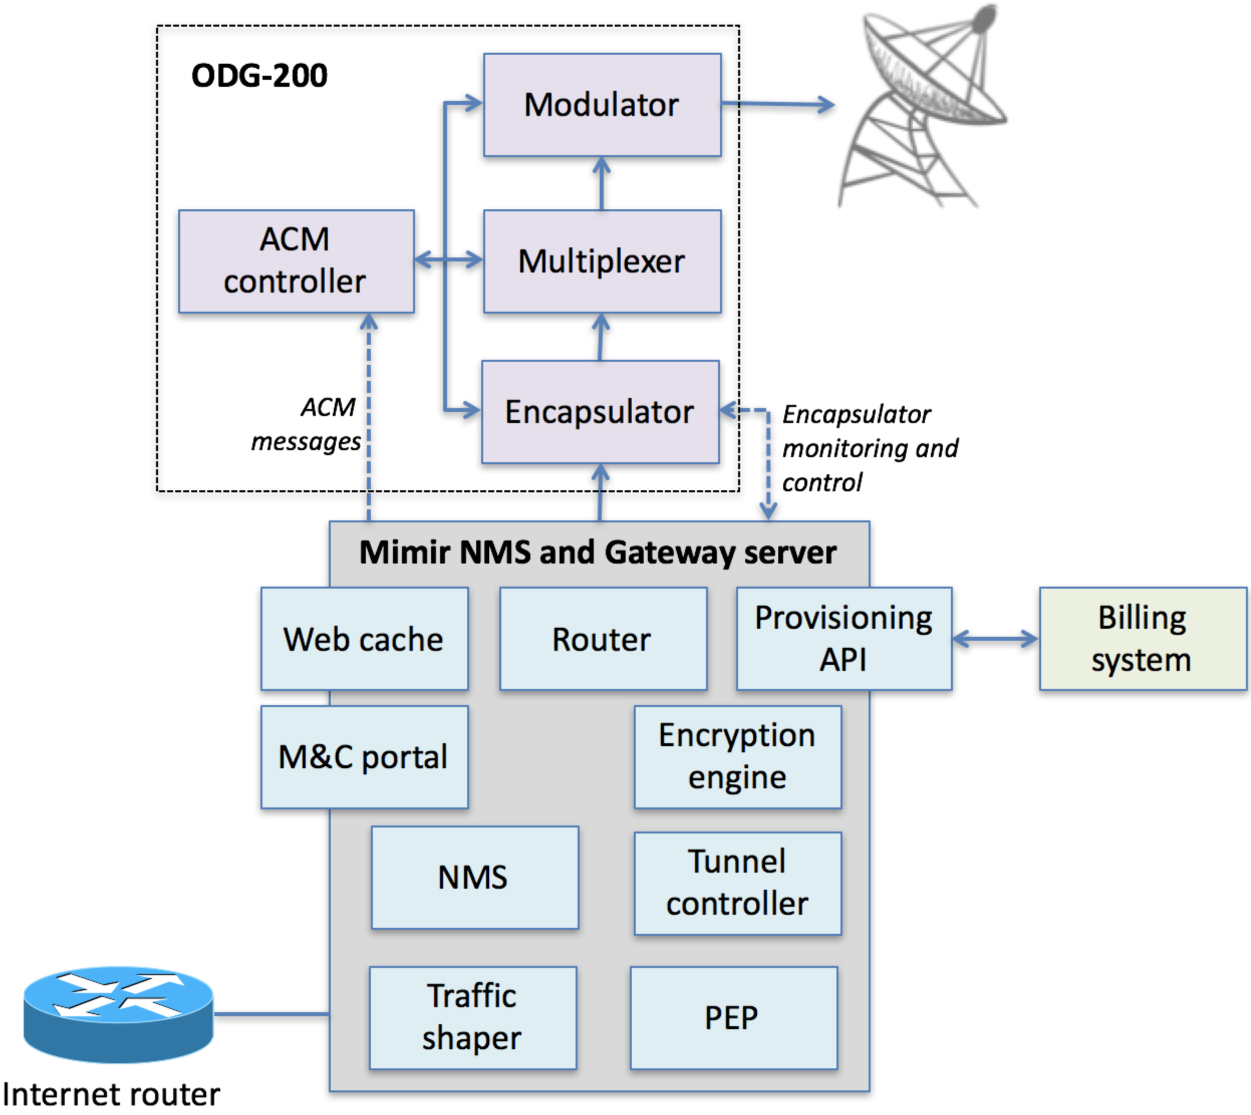

Figure 9. Mimir deployment environment

6.1. Deployment example

Mimir can be used in everything from small installations with a single PID on one transponder to large installations with multiple transponders in several orbital positions.

Figure 8 shows how the most basic Mimir installation would look, with everything running on one server. The example below describes a large installation, illustrated in Figure 11.

The space segment of the installation consists of one satellite with two different beams. One beam has three transponders, the other one transponder. ACM is used on all transponders.

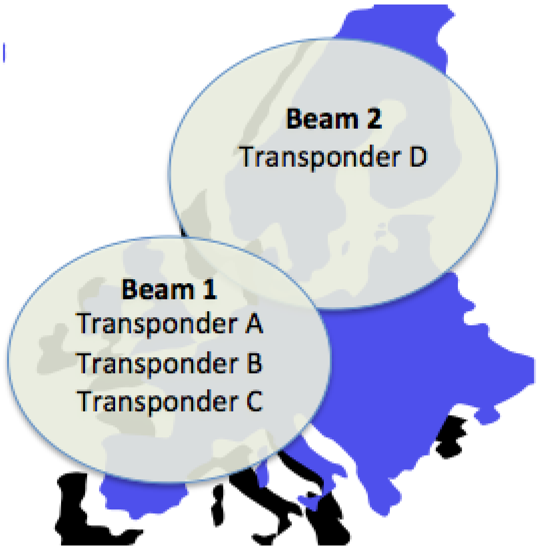

Figure 10. Example coverage map

The system is planned for a large number of users, why Mimir NMS and Mimir Gateways are deployed on separate servers to avoid overloading any component in the system.

The encapsulators are first added to the system. Each encapsulator is defined in Mimir with the satellite parameters required to access its transponder.

In next step three Mimir Gateways are added. One Gateway is configured to be used as backup. Two pairs of primary and backup encapsulator are added to each primary server, all four encapsulators pairs are added to the backup server.

Next two populations are added: Beam_1 and Beam_2.

Four services are created: 1A, 1B, 1C, and 2D

The services are assigned to one encapsulator each.

One FAP and one plan are created for the whole system.

When users are created they are added to one of the two populations depending on their geographical location.

When users in the Beam_1 population connects they will be assigned to one of the services 1A, 1B, or 1C depending on the current service congestion. Users in Beam_2 will always be assigned to the service 2D.

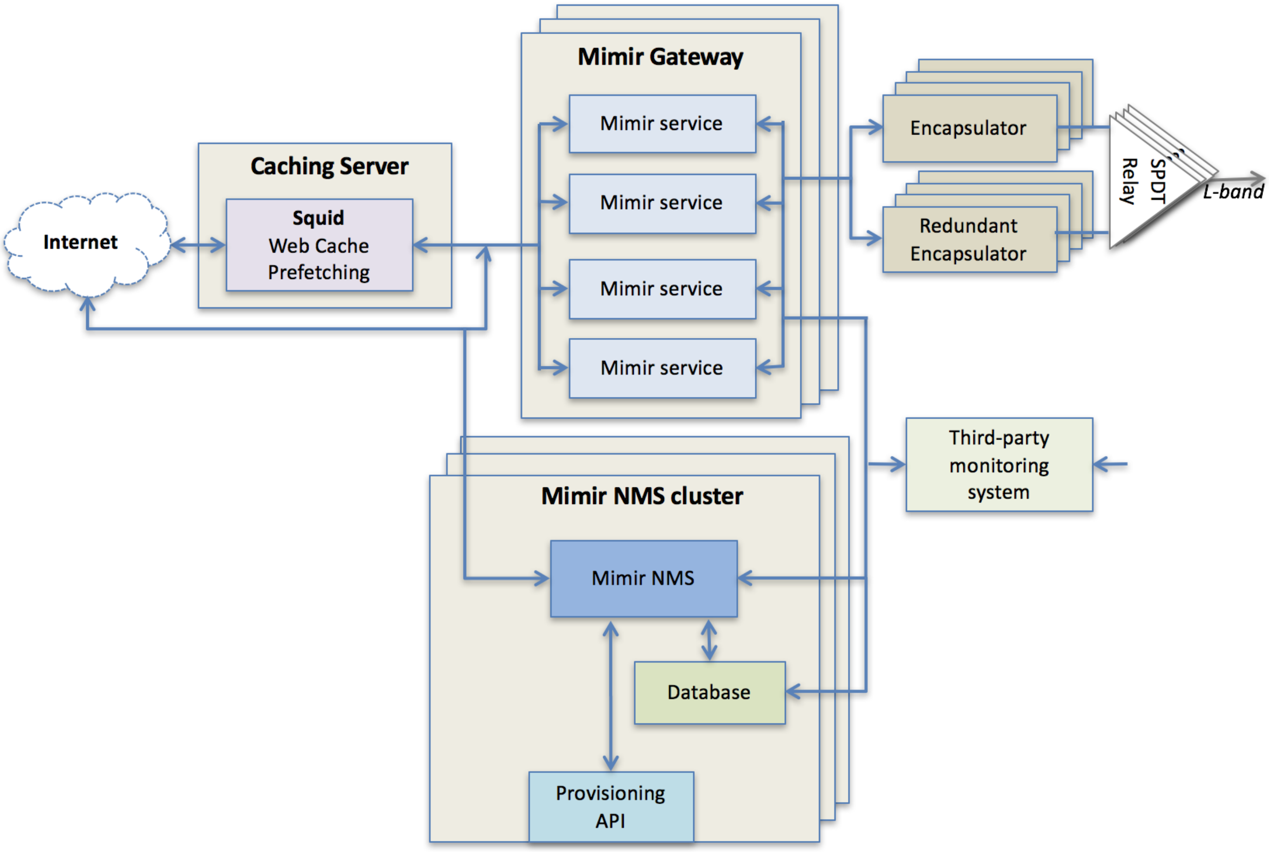

Figure 11. Large installation example

7. Third-party Billing and Provisioning Systems

It is possible to integrate the system with third-party provisioning and billing systems.

7.1. Mimir Provisioning API

The Mimir Provisioning API is an RESTful web service interface for integration with third-party provisioning and billing systems. The API allows third party systems to create and manage users as well as retrieve billing data.