Congestion Free Networks

Technical white paper

Introduction

Chorus operates a congestion free network.

This means that we continuously monitor our fibre and copper networks, and proactively augment capacity before any part of the network fills up. That is, we ensure no links exceed the maximum utilisation threshold, as defined in section 3 of this document.

However, this does not mean that Chorus networks are non-blocking. Chorus networks are designed to statistically multiplex highly stochastic bandwidth demand from large number of customers. Even under lightly loaded conditions, there is a finite probability that you could see some frame loss over very short durations (typically < 1 ms).

Purpose

This whitepaper has been developed to meet the following requirements:

- Provide an unambiguous engineering definition of what a congestion free network is;
- Explain how Chorus achieves this through proactive monitoring and capacity management rules; and
- Show how this applies to different traffic classes.

About this document

This document is a guide to how Chorus manages network bandwidth capacity within our copper and fibre networks and what this means for your services.

Chorus

Chorus endeavours to make this document as comprehensive and technically accurate as possible. However it may need to be updated from time to time to include clarifications, errata or additional content. Feedback on the content, technical accuracy or clarity is welcome and should be forwarded through your account manager.

This document should be read in conjunction with the *Chorus Gigabit Experience* technical white paper available on the Chorus Website.

Congestion free networks

As per the following diagram, the Chorus congestion free network applies to the Access Node and aggregation network, but not to the physical access between the broadband customer and the access node.

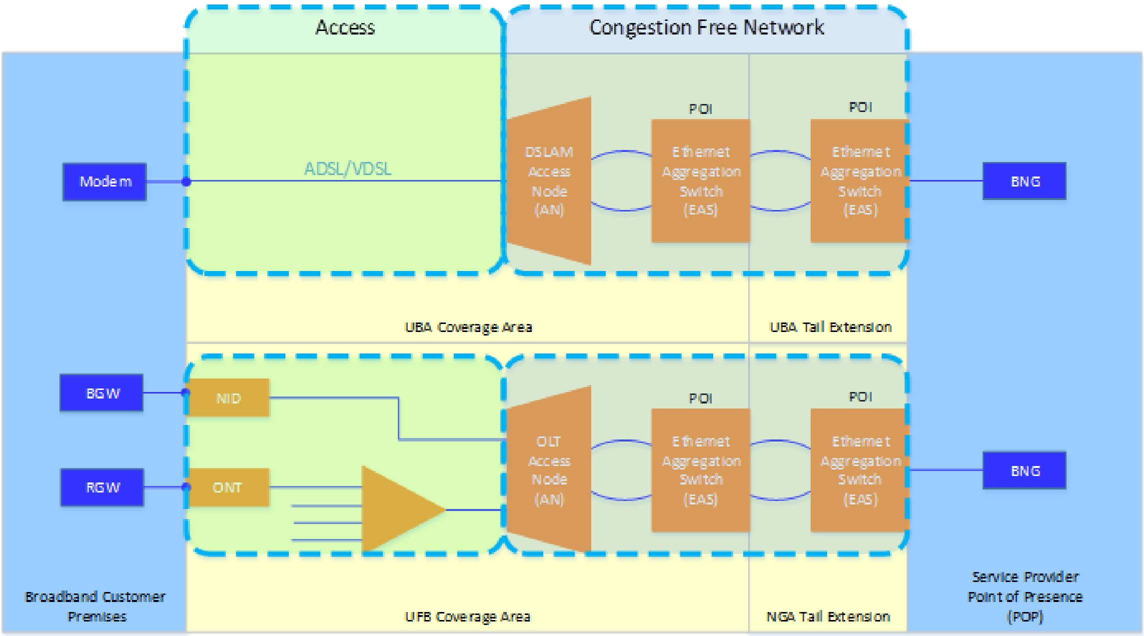

The access lines are excluded due to the low levels of contention of these portions of the network, as described below.

DSL copper access

The DSL copper access is a rate-adaptive point-to-point connection between the customer premises and the DSLAM access node.

The rate-adaptive nature means the available bandwidth per customer varies depending on; distance, line quality and external conditions such as electromagnetic interference or house-wiring.

This is important from a performance perspective, because the access line will get congested if either the customer (upstream) or the BNG (downstream) sends more traffic than the line can support. In general, the downstream path is the more critical direction for customer experience.

For downstream traffic, if frames are sent at a higher rate than the DSL line rate, then the Chorus DSLAM will buffer the excess frames and forward them when bandwidth is available. Weighted Random Early Detection (WRED) policies (see 5.1), are used to manage this overflow, noting that Service Providers can implement their own traffic management policies in the BNG.

For upstream traffic, frames can only be submitted at line-rate. Excess frames are managed within the customer domain and techniques and behaviour may vary, depending on the CPE being used.

GPON fibre access

The GPON fibre access is a 2.4/1.2 Gbps shared fibre access supporting a number of customers:

- Chorus uses a default split ratio of 16:1, which supports an average bandwidth of ~ 150/75 Mbps per ONT;
- Early deployments may use a higher split ratio. These will be monitored and groomed as required.
- Some remote regions (e.g. for RBI) will have lower split ratios.

High speed offers allow customers to get a higher speed than their equal share, by giving them access to any unused bandwidth from other customers on the shared GPON network. The implication is that these customers always have at least 100 Mbps downstream and 50 Mbps upstream available, but bandwidth above this is dependent on the behaviour of other users in the same shared GPON Access, at any given time.

In practice, the stochastic nature of broadband means customers will generally not be constrained - even for higher speed offers. If we look at Chorus' GigCity for example, we can see that the average capacity consumed by a gigabit broadband customer is much less than their peak capacity.

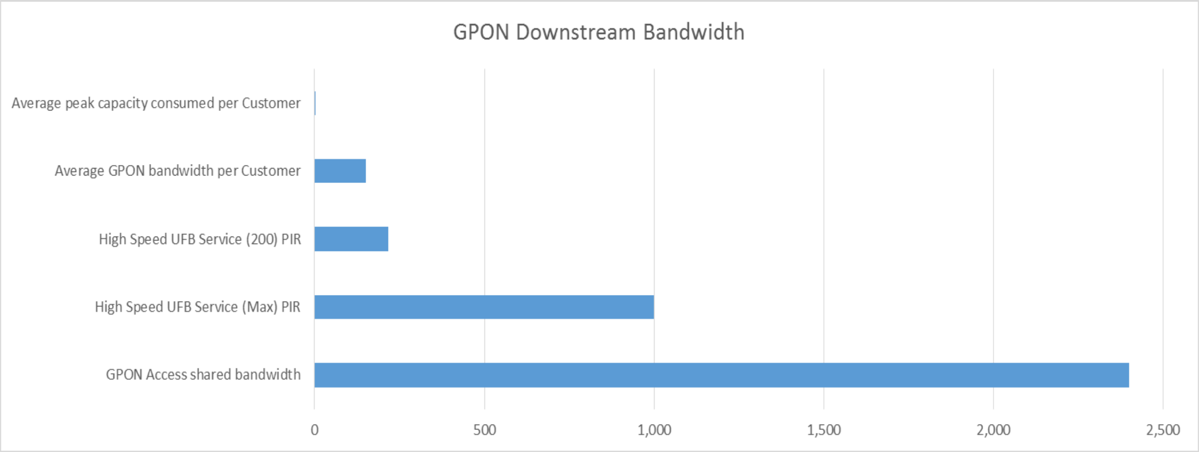

This graph demonstrates the key benefits of contention. The highly stochastic nature of internet services means that an individual customer can experience speeds much higher than the average bandwidth available and it would be very rare for any gigabit customer to not get the bandwidth they are asking for.

Point to point fibre access

Like DSL, point to point fibre accesses use a dedicated connection per NID, per Customer. However services are policed on ingress, so excess frames (greater than service rate) are immediately discarded rather than queued.

With NGA Business Premium (currently under development) it is possible to configure multiple services on a single access, in which case, it's possible that the aggregate bandwidth of these services exceeds the access speed of the line.

However under single service scenarios, the Chorus congestion free network policies apply to all frames within the service profile.

Constrained DSLAMs

Approximately 2% of broadband customers are on legacy technologies. Typically in remote areas, their limited backhaul can result in congestion, particularly in peak demand times.

Chorus has an active programme to replace these legacy technologies, which will move these customers to a congestion free network. This programme is a combination of standard lifecycle replacements, co-ventures with local communities or as part of the Rural Broadband Initiative (RBI).

Definition of a congestion free network

Chorus defines a congestion free network as one where all Internal Network to Network Interfaces (I-NNIs) average bandwidth, as measured over [15] minutes, is less than [95%] of the I-NNI physical line rate.

Where:

- [15] minutes is the time interval over which network utilisation is measured;
- [95%] is the maximum utilisation threshold, where the network is considered to be non-performing.

This includes:

- I-NNIs between Access and Aggregation domains; and
- I-NNIs within the Aggregation domains.

It excludes:

- Handover Connections (E-NNIs). Contention of these ports is controlled and managed by the service provider
- Access lines (DSL ports, PON ports. These are low aggregation points (1:1 or 16:1).

The I-NNI average bandwidth is measured in both directions, using the following formula:

$$Average\ Bandwidth = \frac{Total\ Octets\ delivered\ over\ [15] minutes}{[15]\ \times\ 60\ seconds}$$

The Total Octets are as per RFC2863 (ifInOctets/ifOutOctets), which includes framing characters but excludes Ethernet preamble, start frame delimiter (SFD) and interpacket gaps.

The I-NNI port utilisation is then calculated based on the following formula, expressed as a percentage:

$$Port\ Utilisation = \frac{Max(\ Upstream\ Average\ Bandwidth, Downstream\ Average\ Bandwidth)}{Port\ Physical\ line\ rate}\ \%$$

This definition takes advantage of the smoothing effect on high aggregation paths, which reduces the variability of bandwidth demand within the measurement interval. Chorus uses 15 minutes as a port utilisation measurement interval. Testing shows this provides an accurate indication of port utilisation while balancing the collection and processing overheads. Comparison tests with 5 minute collection intervals showed less than 2-4% difference in results between 5 and 15 minutes collection and any negative variance would be picked up by calibration nodes.

Although both upstream and downstream average bandwidths are calculated, only the maximum value is used for port utilisation as this will drive any decision to augment capacity.

The following graphs are examples of bi-directional 15 minute collection intervals on a 10 Gbps fibre OLT-aggregation I-NNI over a 24 hour period:

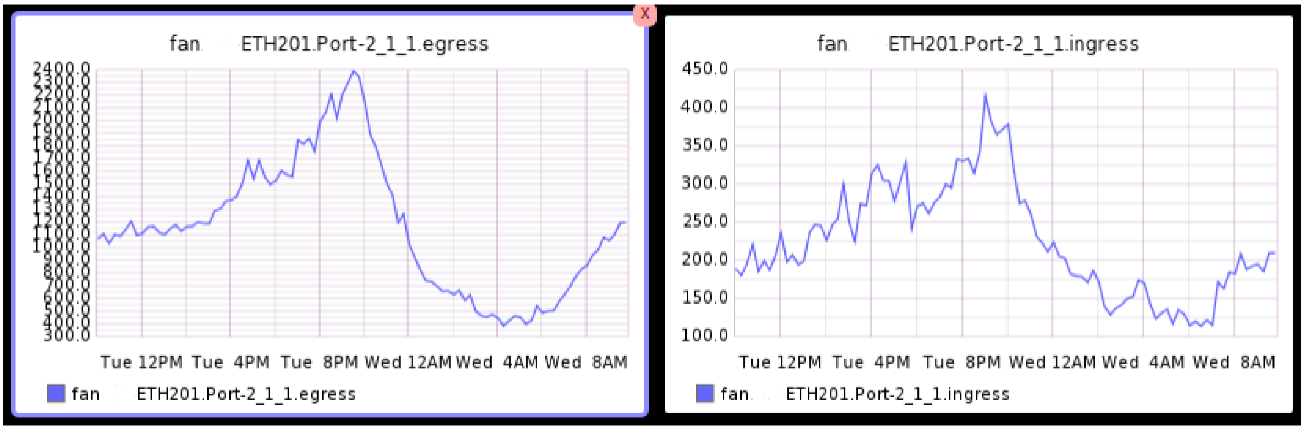

It is not appropriate to use this definition on contended low aggregation points, such as (EIR) over-subscribed PONs, as the high ratio between PIR and average bandwidth will result in high variability of bandwidth within the measurement interval.

Chorus capacity planning policy

The maximum utilisation threshold is a key planning parameter that drives Chorus' capacity planning model as follows:

- The maximum threshold should be a failsafe figure, i.e. broadband customers **should never experience** the network at this level, under normal operating conditions.
- Chorus capacity planning rules are therefore designed to ensure that this figure is never reached, except in extreme circumstances.

The following graph shows how Chorus capacity planning works, using default thresholds for fibre:

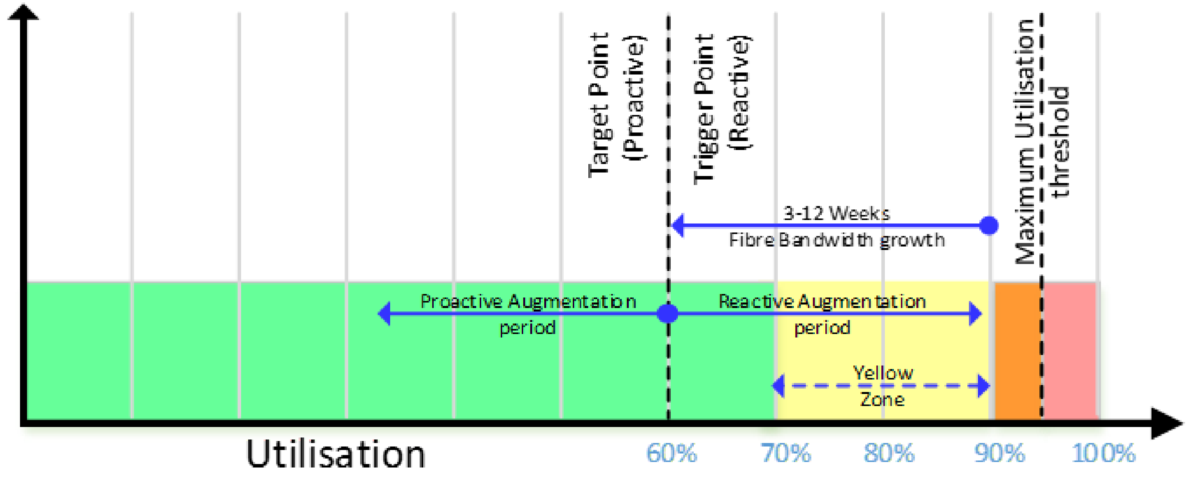

Zone	%	Congestion	Experience	Description
Green	0-70%	Low	Good	User experience is good. Investment is only done to prevent utilisation reaching the yellow zone. This is the normal network condition.
Yellow	70-90%	Low	Good	User experience is good, but investment is underway to reduce utilisation to the target point threshold. Capacity planning aims to avoid this region as much as possible. If the network enters, or is predicted to enter, this area then capacity investment will be expedited to return the network to green as soon as possible.
Amber	90-95%	Low-Medium	Med	Customers are still getting a fast service but performance degradation may start to be noticed. The network is never intended to enter this space, except under extreme conditions, such as abnormal traffic conditions or massive unpredictable growth.
Red	95%+	Med-High	Med-Poor	Customer experience will deteriorate rapidly. This is the failing zone.

Where:

Term	Description
Target Point	This is the <u>two-year forward-looking</u> capacity design target, based on a revolving two-year view of predicted traffic. The forecasts are redone quarterly. This is Chorus' best estimates of bandwidth demand over the next two years based on a combination of historical and current trends, market research on bandwidth growth influencers and international trends. These trends and predictors are reassessed each quarter.
Trigger Point	This is a reactive trigger, looked at weekly, to take into account unforecasted high bandwidth growth occurring between quarterly reviews. If the network reaches this point, Chorus investigates the risk of the network entering the yellow zone prior to the planned investment cycle and adjusts capacity investment appropriately. It is not necessarily a trigger to immediately invest. The investment decision will be based on; avoiding the yellow zone, likely bandwidth growth for that link, and noting that some links are considered sensitive and some non-sensitive (see definition below).
Yellow Zone	The yellow zone is where we consider broadband customer experience satisfactory but there is a risk of performance degradation if the network capacity is not augmented before bandwidth growth result in port utilisation exceeding this zone. Capacity planning is designed to prevent the network entering this zone. However unpredictable traffic growth (i.e. within a short period of time) could result in this zone being entered. The width of the zone is a direct function of growth and time: Fibre is currently seeing a high growth of per user bandwidth, plus a growth in user connections. This results in an average bandwidth growth of ~ 200% per annum. This means a 20% yellow zone can be traversed in less than 8 weeks under normal growth. Copper would support a narrower yellow zone due to lower growth. This is also offset by copper to fibre migration. Note that this growth is uneven and the zone needs to be wide enough to cope with worst-case areas. The yellow zone is not considered a failure condition, but requires network augmentation to prevent the network entering the red zone.
Maximum Threshold	The maximum threshold is the point at which the network is considered failing, with a buffer between this threshold and the yellow zone. This parameter drives Chorus' internal capacity planning thresholds and business rules, intended to prevent the threshold from being exceeded. Setting this figure too low will require a combination of tighter business rules (which reduces service providers' flexibility), and force Chorus investment in over-capacity. This risks stranded assets as customers migrate from copper to fibre.

Predicted BW growth	Chorus undertakes long-term bandwidth growth analysis. It looks at current trends in; growth per user, uptake, and considers local, regional, national and international patterns in bandwidth growth. This is our best estimate of where we expect bandwidth demand to be in two years.
Unpredictable BW growth	Unpredictable growth is where bandwidth grows quickly, within a quarter. It could be triggered by a number of different factors such as: The introduction of new high bandwidth content, like Netflix or a similar streaming service. A change in broadband customer behaviour, such as a sudden trend from HD to UHD video, which could be driven by content provider pricing; Service provider network grooming. **Such growth will not necessarily impact the network evenly.**

It should be noted that growth rates vary considerably; across the network, between copper and fibre, connection growth, migration from copper to fibre, and user bandwidth. Thus the width of the yellow and amber zones will vary link by link.

This is true even within the high growth fibre network. For example, when a new OLT is installed the bandwidth consumption will be dominated by growth in connections, but as the OLT fills up the growth in average bandwidth per user will become more important.

Implications on traffic classes

The congestion free network has different implications for different classes.

Best efforts/low traffic class

A congestion free network does not guarantee zero frame loss, even in the green zone. Frame loss is a function of discard eligibility, buffer occupancy and Weighted Random Early Detection (WRED) slopes.

Chorus uses WRED slopes for congestion avoidance. This is a much fairer approach than a tail drop algorithm and provides a much better broadband customer experience. We consider this best engineering practice and is in place to ensure smooth TCP performance during (short duration) congestion events, which can occur even under low average traffic demand conditions.

The following graph shows the current Chorus WRED policy for fibre services. Given the changing nature of traffic patterns, Chorus regularly reviews this policy and welcomes industry feedback to ensure optimum configuration. However, Chorus strongly advises against turning WRED off, due to the impact on TCP traffic.

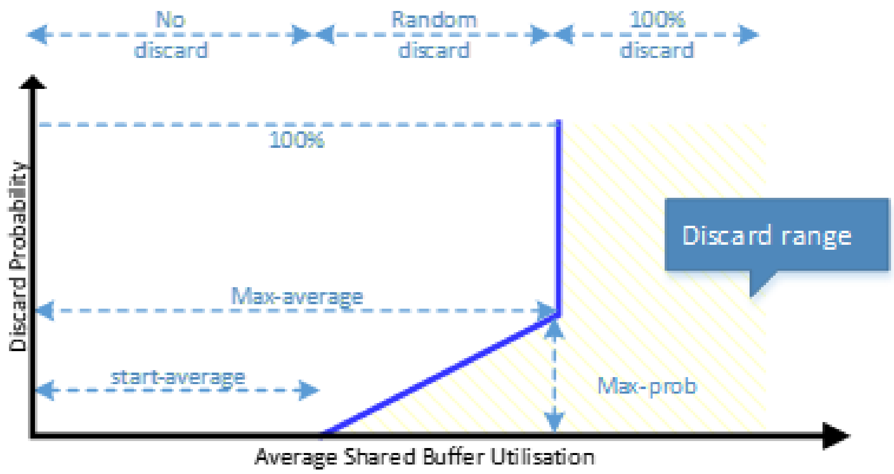

Buffer Pool	Start-Average	Max-Average	Max Probability
Network Egress	50%	75%	80%

WRED is based on buffer occupancy, which is based on traffic demand on a particular link over milliseconds, rather than link utilisation, which is measured over minutes. There is a finite probability, even at low average link utilisations, that a sudden burst of frames on a particular path will increase the queue occupancy such that WRED is triggered.

This increase in probability is exponential, as shown in the table below.

Zone	Green	Yellow	Red	Red
Utilisation	0-70%	70-90%	90-95%	95%+
WRED probability	Very low	Low	Low-med	Med-high

The 100% discard is applied at 80% buffer utilisation to allow for WRED processing time.

The following graph shows an example of this policy overlaid with the different capacity planning zones. This graph does not show a normal link experience, but does show the relationship between WRED frame loss and utilisation.

As can be seen, there is a small probability that frames will be WRED discarded even at 60% average bandwidth, and this probability exponentially increases, particularly after 95%.

This graph shows that frame loss is low (~ 0.1-0.2%) even at the maximum utilisation threshold of 95% port utilisation. Given the maximum utilisation threshold is a failsafe that should never be achieved under normal operating conditions, Chorus believes 95% is a good balance that drives the right capacity investment behaviour without constraining service provider flexibility or degrading broadband customers' experience.

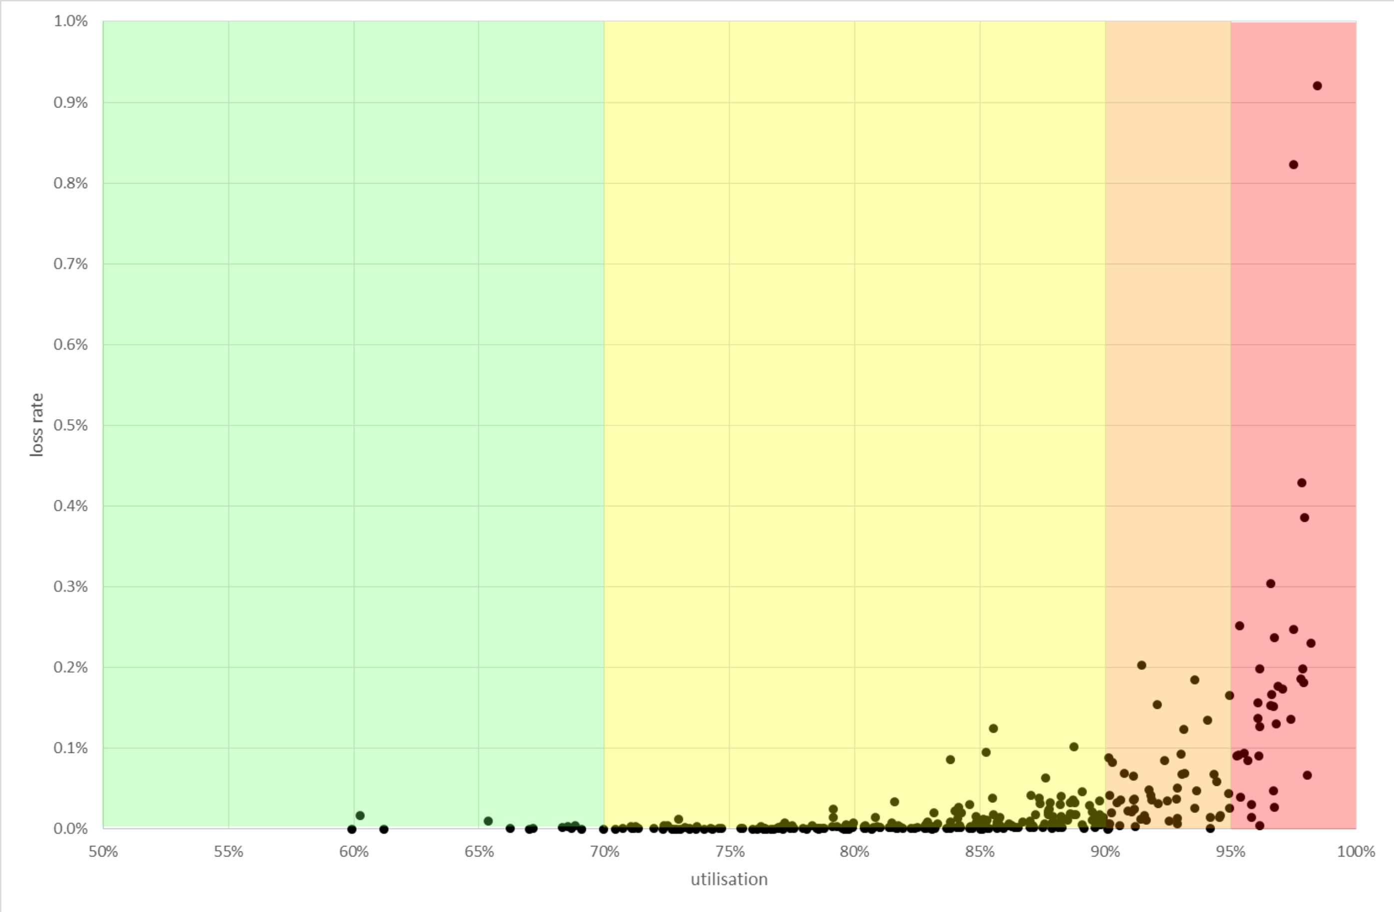

The impact of frame loss on typical internet (TCP/IP) traffic is the potential retransmission of lost frames, which ultimately slows the potential maximum observed throughput. The impact of this will vary as many applications and content servers use a variety of techniques to mitigate this effect, including multiple (parallel) TCP sessions, modern TCP rate control algorithms, window scaling and support for selective acknowledgements.

Real time/high traffic class

Although WRED policies are applied to CBS, and thus CIR frames, the QoS policies are specifically designed to avoid High traffic class frame discards.

The High traffic class has a separate WRED policy to the Low traffic class. We do not expect frames marked as High to be discarded, as the volume of Low frames far exceeds the High frames, and High frames are strictly prioritised. WRED is still assigned to ensure traffic degrades gracefully, even under extreme conditions.

The following graph is a snapshot of an I-NNI link that shows the different volumes in Low (Blue) and High (Red) frames:

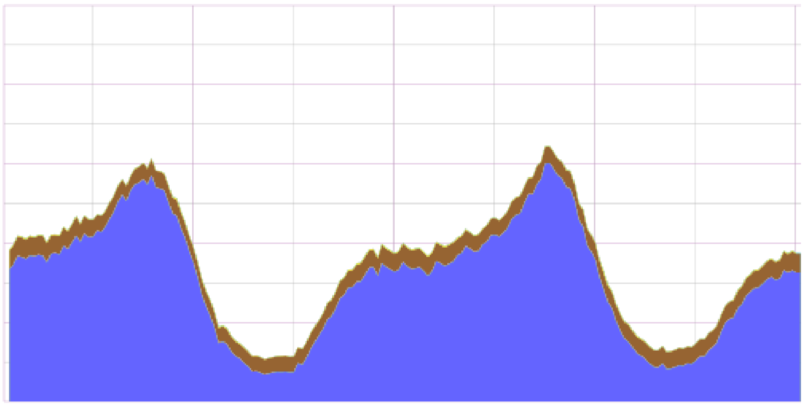

It should be noted that High traffic class demand is relatively smooth throughout the day, while Low traffic class demand increases until the (current) peak hour of around 9pm. Therefore investment in broadband capacity in New Zealand is driven by peak hour, Low traffic class.