DEIP
DISTRIBUTED ENERGY INTEGRATION PROGRAM

CSIP-AUS EXPLAINER

May 2025

Aligned with CSIP-AUS version 1.2

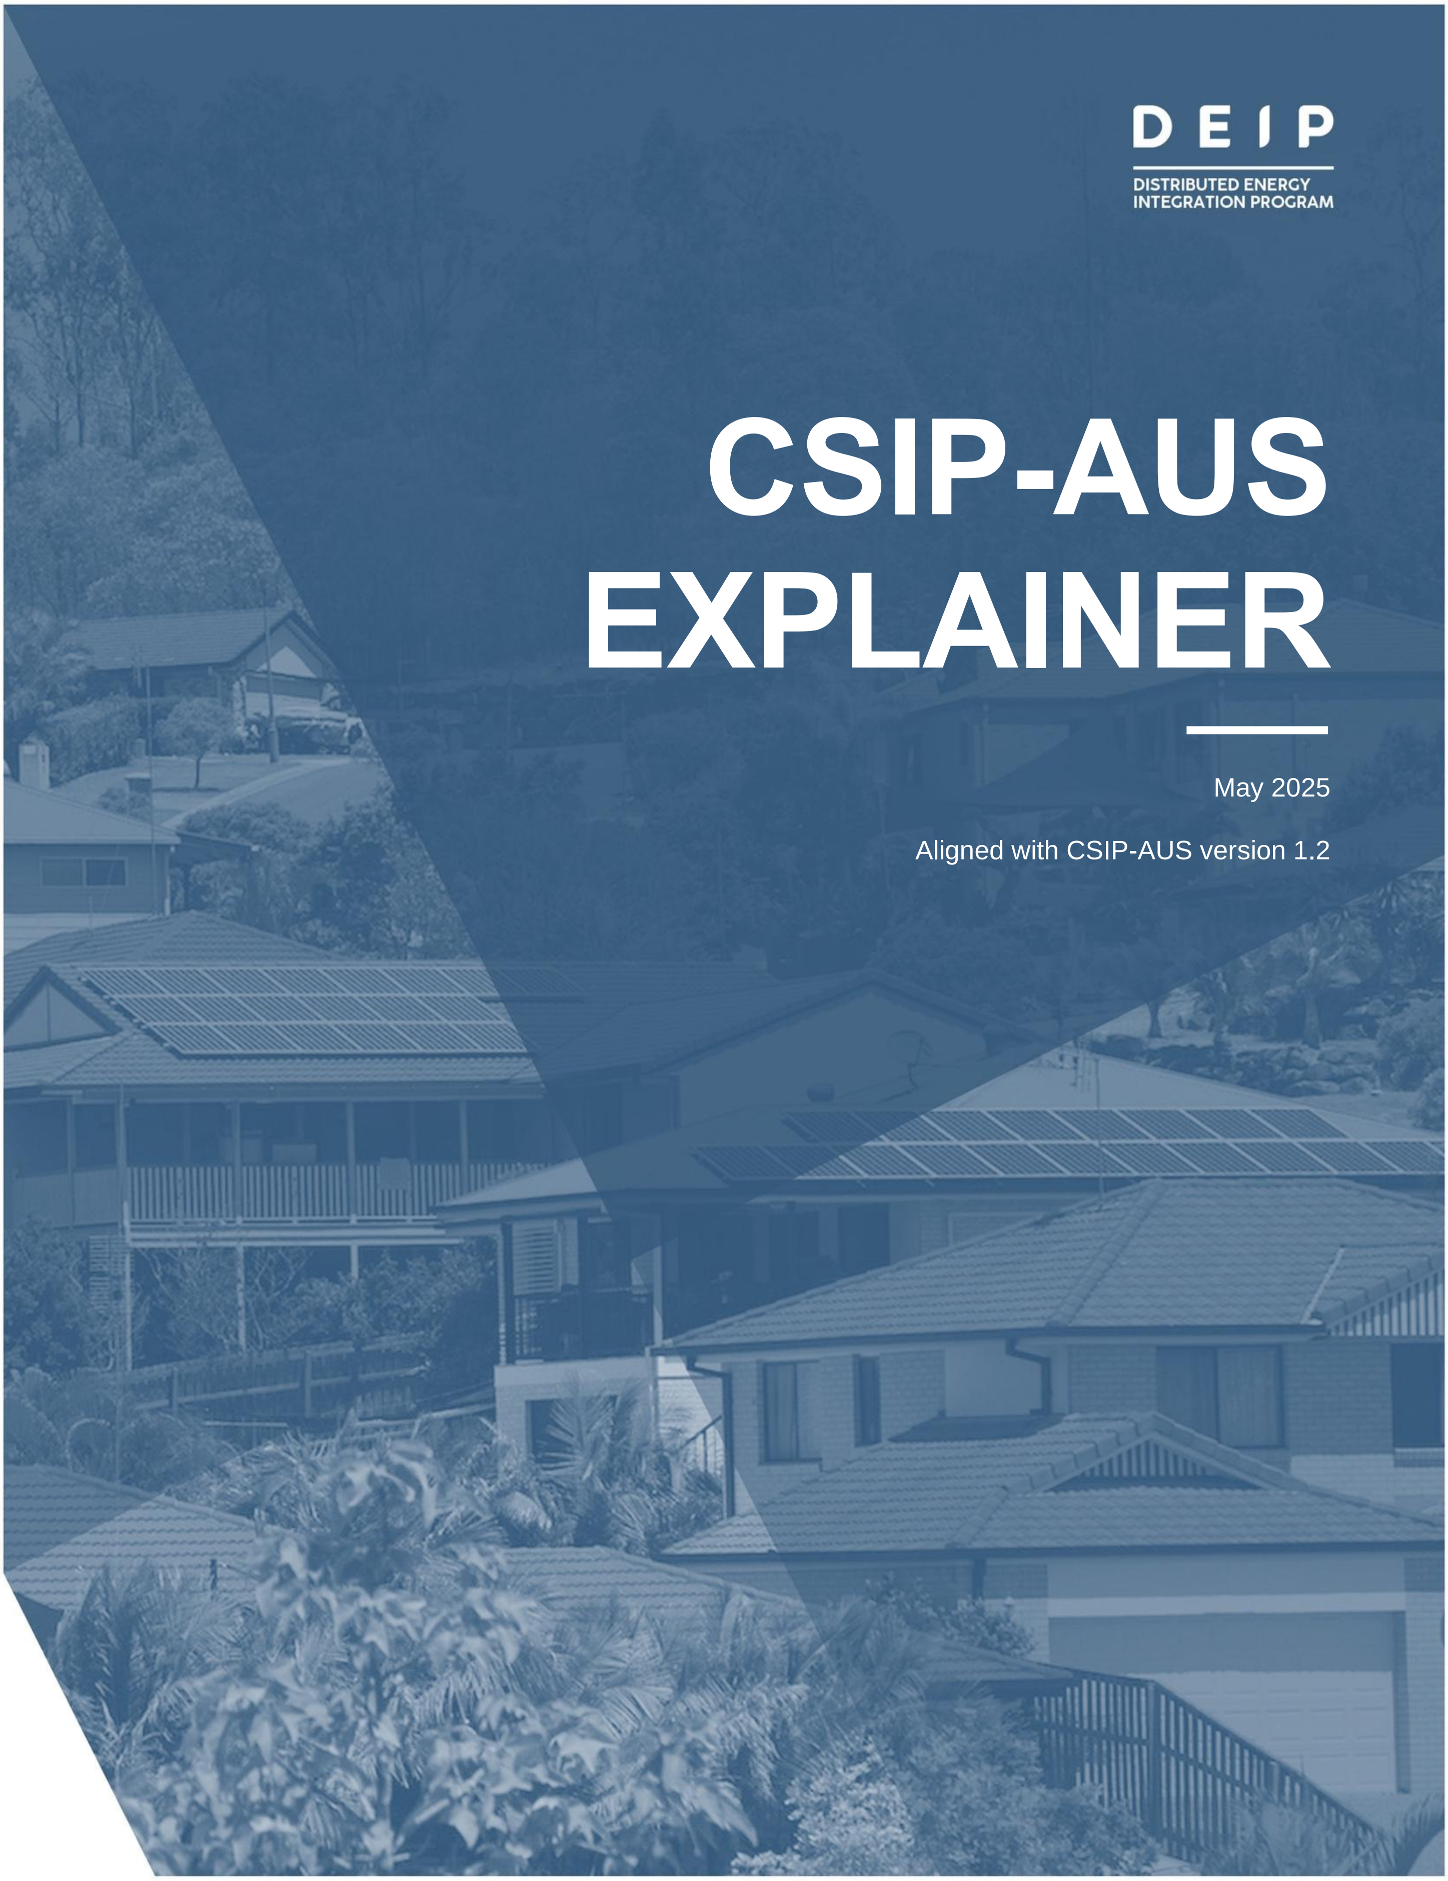

ABOUT DEIP

The Distributed Energy Integration Program (DEIP) is a collaborative network of government agencies, market bodies, peak industry bodies and consumer associations working together to maximise the value of distributed energy resources (DER) for all Australian consumers.

The DEIP members exchange insights, seek industry consensus, and focus attention on priority activities to provide the necessary pre-policy evidence needed to support informed decision making in the Australian energy transition.

ACKNOWLEDGEMENT AND DISCLAIMER

This ‘Common Smart Inverter Profile – Australia’ and this accompanying Explainer document were developed by the DER Integration API Technical Working Group. This working group formed in 2019 as a collaboration of Australian energy sector businesses from across the supply chain, including numerous distribution networks, retailers, equipment manufacturers and aggregators.

CONTENTS

HOW TO USE THIS DOCUMENT

This document is intended to serve industry as a means of improving understanding of CSIP-AUS and its technical implementation.

Sections 1-3 are intended to support readers that are not directly involved in CSIP-AUS integrations, but who may have a desire to better understand the role and high-level functions of CSIP-AUS.

Sections 4-10 are intended for readers with more substantive involvement in integrations, and as such has a much more technical focus.

1 GLOSSARY

The terms below are those that have emerged organically through the implementation of CSIP-AUS across Australia.

Term	Definition
Capability test	The test that is conducted on-site post installation and device registration of the inverter or DER. The purpose of the capability test is to: - Confirm server connection - Validate telemetry is accurate and at required posting intervals. E.g. Current Transformers (CT's) are installed in the right direction. - That site responds to controls at required polling intervals
CEC SCC listing	The Clean Energy Council maintain a listing of those inverters that have passed SAPNs testing. This listing is also updated on the CEC API whereby inverters have the tested SCC as child records under the inverter.
Client	A communications client that exchanges CSIP-AUS control and telemetry signals with a utility server, and enacts them on one or more DER. Used interchangeably with, and equivalent to, the terms *DER client*, *utility client* and *communications client*.
CSIP-AUS Aggregator	A CSIP-AUS aggregator is not the same as a Virtual Power Plant (VPP). Historically in Australia the term VPP and aggregator have been used synonymously, as an aggregator is seen as an aggregation of energy resources. However, with the introduction of CSIP-AUS and the language used in CSIP, an aggregator simply means the cloud platform of a service provider aggregating the IEEE2030.5 communications for end devices. There is no correlation with aggregating energy. This will change in the future as IEEE2030.5:2023 is defining the communications aggregator as a proxy to remove the confusion.
Discovery	The process of a client establishing communications with a utility server and exploring initial resources to gather necessary end-point links to enable ongoing communications.
In-band registration	In-band registration is the process whereby an LFDI and it's associated connection point reference (in most instances will be the site NMI) is automatically inputted into a utility server via the IEEE2030.5 protocol through use of the CSIP-AUS connection point extensions. With in-band registration the LFDI is often not seen by users.

Inverter in-built software communications client	Inverter in-built as used in the CEC listing means that the inverter does not rely on a third party or gateway device for the CSIP-AUS integration. Inverter in-built could be via a cloud aggregator or direct device integration.
LFDI	Long Form Device Identifier
OEM Onboarding or Onboarding	The process of: 1. Validating a client and server interoperability using test certificates. In most cases on test systems. 2. Once validation is completed, the issuing of PKI certs for production. 3. Updating of listing and installer/customer facing systems reflecting on boarded status 4. Out of band registration of aggregator in production systems
Out of band registration	Out of band registration is where the LFDI is entered into the utility server either manually, or automatically through a back end process. This involves the installer or OEM providing the LFDI to the utility via an agreed process.
Static Operating Envelope (SOE)	For the emergency backstop use case where a site is not under a flexible or dynamic connection agreement, the export limit that is being sent via CSIP-AUS is up to the approved export limit. This is referred to as the Static Operating Envelope. Not to be confused with project Converge that creates “Shaped Operating Envelopes”
Utility	Operator of the SEP2 utility server, which is often a network operator or retailer business.

2 INTRODUCTION TO CSIP-AUS

2.1 The Purpose of CSIP-AUS

The continuing expansion of distributed energy resources (DER) in Australian electricity networks is an incredible success story. This is especially true for distributed solar energy resources, with 1 in 3 Australian households now having rooftop solar. But it also presents a significant challenge for traditional network operation, as electricity generation and load in the network is now much more dynamic. Good visibility and control of this DER is required to keep networks stable. This control is also required for customers and other industry participants to make the most effective use of their DER assets.

The *Common Smart Inverter Profile – Australia* (CSIP-AUS) is a communications protocol that facilitates visibility and control of DER in Australia. Network operators are requiring new DER devices to conform to this protocol, so that network stability can be protected, and resources can be more efficiently utilised. Although the timing and circumstances vary from one jurisdiction to another, the intention is to see CSIP-AUS adopted nationwide for all DER over time.

The key use cases enabled by CSIP-AUS are covered in detail in Section 3, but they include:

- *Emergency Backstop*: to ensure network stability during periods of very high solar generation.
- *Flexible Exports*: to allow or limit the total amount of power that can be exported by a connection to the network.
- *Flexible Demand Reduction*: to reduce consumption by devices (e.g. battery chargers) that can vary the amount of power they draw.
- *Generation Controls for Microgrids*: which mixes controls in Emergency Backstop and DOEs as these apply to microgrids.

These, along with other use cases, are essential to manage efficient and safe operation of DER in the Australian energy network.

2.2 Key Concepts

Here are some of the fundamental concepts that CSIP-AUS relies upon.

Connection model: CSIP-AUS connects DER devices (solar inverters, batteries, EV chargers, etc) to Utility Servers belonging to network operators (typically, Distribution Network Service Providers – DNSPs) or retailers. The Utility Server can send control instructions to DER at each site and get responses from that DER to verify its current operating state.

Connections can be **direct** to the DER device, where the device implements the client side of the CSIP-AUS protocol. Or the device might use a cloud-based **proxy** connection (sometimes called an "aggregator" connection), where the cloud-based proxy implements the required CSIP-AUS client protocol. The proxy translates this to proprietary communications with the DER to achieve the required behaviour. In both cases, the communication is typically over the Internet and leverages the customer's internet connection for the end communications.

Connection can also be to a site-based **gateway** device. In this case the gateway device implements the CSIP-AUS protocol and communicates to DER devices on the site using proprietary protocols or (CSIP-AUS). Connection from the gateway to the Utility Server is typically then also over the Internet and can be direct to site or via the gateway service providers cloud-based proxy connection.

Figure 1 below shows examples of these different connections for some networks in Victoria.

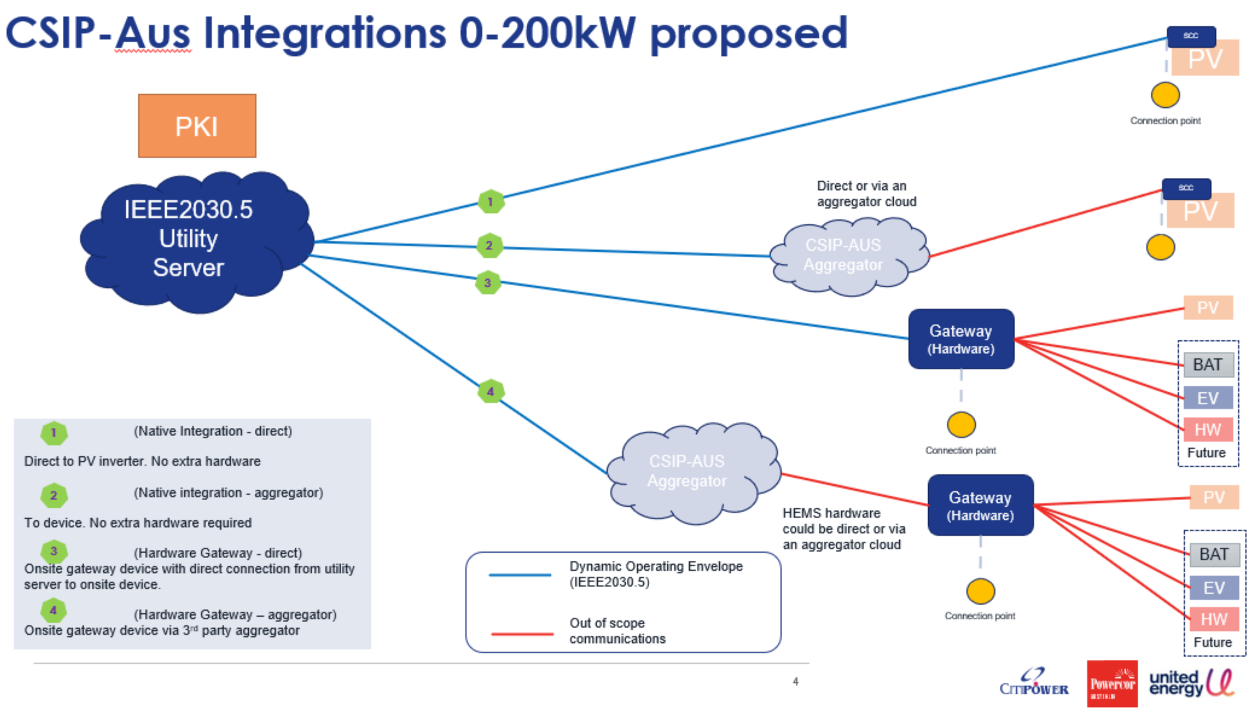

Figure 1: Connection model implementation examples from Victoria

Scope of control: Most CSIP-AUS controls need to operate at a **site** level, ensuring that overall limits for multiple devices operating on a site are enforced. Limits tend to apply at a point of common coupling (PCC), which is typically the metering point for a customer's site. Individual assets are thus often referred to as "behind-the-meter" (BTM) assets. Where site limits need to be enforced, CSIP-AUS expects that there is some client device capable of co-ordinating the behaviour of the on-site assets to ensure that overall obligations are met.

However, CSIP-AUS can be used to monitor and control **individual** devices. For example, a single solar inverter might be instructed to not export any power to the grid for a period of time. For some sites, control of an individual device will effectively control the site.

Dynamic Operating Envelopes (DOEs): These are flexible limits on how much electricity a site can export to or import from the grid at any time. Unlike fixed (or static) limits, DOEs adjust based on real-time grid conditions, and can control DER behaviour to assist with safe and efficient network operation.

A key goal of DOEs[1] is to maximise customers' ability to export their solar, while providing a mechanism for network operators and AEMO to manage local and system constraints. In addition to benefiting the customer receiving the DOE, this technology also reduces costs for all network and system users, through increased network utilisation and displacement of other forms of generation. As explained in Section 3, DOEs can be used for peak load management ('flexible imports') to curtail electricity demand, or to curtail generation during system-wide emergencies ('emergency backstop').

CSIP-AUS DOEs apply at a site level.

[1] https://arena.gov.au/assets/2022/03/dynamic-operating-envelope-working-group-outcomes-report.pdf

Import and Export Limits: The DOE functionality in the CSIP-AUS enables the communication of range of DER management signals including import and export limits at a site level. Depending on the use case, it may include one or both of these limits. The most common application is for communicating export limits for rooftop solar.

An *export* limit specifies the maximum power exported to the grid from a site. It does not restrict generation that is used by a customer at their premises. This allows DER system owners to continue to realise value from production of their system through self-consumption while ensuring physical constraints of a local network are adhered to.

An *import* limit specifies the maximum power imported from the grid to a site to drive flexible loads (e.g. batteries and EV chargers). The intention is to allow customer to access faster charging speeds when the network capacity is unconstrained, while providing a mechanism for load reduction at times of high network loading that threatens network security. CSIP-AUS import limits do not apply to uncontrollable loads, although these may have other limits that apply (e.g. current limiting through circuit breakers and/or fuses).

Generation and Load Limits: In addition to export and import limits that apply to energy supplied to or drawn from the grid, CSIP-AUS also allows management of total controllable generation or load on a site. These controls apply to the net energy generated or consumed by all devices that are managed by the CSIP-AUS client on a site, however they do not need to respond to generation or load from unmanaged devices – this is detailed further in the example below.

A *generation* limit specifies the maximum net power that managed DER on the site can generate. The intention is to be more restrictive than an export limit, as local unmanaged load would not otherwise allow the site to generate higher power while still meeting the limit, and would typically be used when low export limits have not proved sufficient to mitigate minimum load issues that could risk destabilising the energy network or system.

A *load* limit similarly specifies the maximum net power that managed DER on the site can consume. As with generation limits is intended to be more restrictive than an equivalent import limit, and would typically be used in a similar manner – to enable stronger limits to be applied when import limits have not mitigated the risk of a peak load event impacting network operation or customers' supply.

Example **-** consider a household includes a battery (B) and an AC-coupled solar system (S1) that was being managed by a CSIP-AUS client, and also a second AC-coupled solar system (S2) that is not managed by the CSIP-AUS client. If the CSIP-AUS client receives a generation limit of 3kW, this indicates that the managed devices B & S1 will not collectively generation more than this 3kW limit. As this is a net limit, if battery B is charging – say at 1kW – then the solar system S1 can generate up to 4kW without them collectively generating more than the 3kW limit, and so this limit is still being met. Additionally, as S2 is not a managed by the CSIP-AUS client, any generation it provides is not considered as part of this limit, and hence it can generate unconstrained. Similar principles would apply to unmanaged load at the site – for example, the residual household load does not need to be considered when determining conformance with load limits.

De-energise and disconnection controls: De-energise and disconnection signals provide mechanisms for a utility server to shut down the operation of all managed DER on a site.

A *de-energise* signal instructs the client to force all managed DER to no longer generate or consumer any energy. Unlike generation and load limits this means that all managed DER would individually cease generating and consuming (rather than allowing netting of energy within the collective of devices). This would most likely be used during a significant grid event – for example, an extreme minimum system load event - to force all generators or loads to cease producing or consuming power. Utility servers using this control should be cognizant that is applies to both generation and load, and ensure that by applying this they don't inadvertently make the problem worse by *increasing* net load (or generation).

A *disconnect* signal instructs the client to electrically disconnect all managed DER form the grid supply. For inverters confirming to AS/NZS 4777.2 this would typically be done through the operation of the automatic disconnection device. This control can be used when electrical connection itself provides a risk to the power system, for example during a cyber-attack where DER may be used to impact grid operation. As with de-energise commands this will cause the devices to cease both generating and consuming power, and carries similar risks as described above.

Managed and unmanaged devices: a customer site may include both CSIP-AUS managed devices and those which cannot be (or are not required to be) managed by CSIP-AUS. The arrangements for such sites are complex and depend on a range of factors (government regulations, desire to make best use of pre-CSIP-AUS investments, feasibility of upgrades, etc), so care must be taken to work through the expectations for each use case.

In some cases, the requirement may be that unmanaged devices should conform to static limits, which are typically stricter than the dynamic ones, while managed devices have more flexibility. But this won't be the case for all use cases in all jurisdictions.

Fallback behaviour: CSIP-AUS also specifies fallback behaviour that applies when communication between the DER and the Utility Server is lost (e.g., during a telecommunications outage). This means that customer premises must have local control capability to conform to specified fall-back behaviour, including sites with multiple DER devices, which may require local orchestration to meet fallback requirements.

2.3 CSIP-AUS, IEEE 2030.5 and CSIP

CSIP-AUS is an adaption of the CSIP implementation guide for the international smart grid communications protocol standard IEEE 2030.5. CSIP-AUS establishes how the IEEE 2030.5 protocol should be used in Australia to control DER. The aim of CSIP-AUS is to precisely define how DER should be controlled so that capabilities (such as DOEs) and use cases (such as emergency backstop) can work consistently in Australia. The main extensions introduced in CSIP-AUS are:

- To recognise export and imports at a site, thus enabling self consumption. This is why CSIP-AUS site require an export meter. Without an export meter, controls would apply to the DER directly and therefore impact customer self consumption.
- Connection point extension for a prescriptive way for in-band registration, thus enabling large scale registration of DER to utility servers in a standardised way.

The vast majority of the functionality defined in CSIP-AUS is specified within IEEE 2030.5. However, CSIP-AUS does include a small number of extensions to the protocol itself. Moreover, IEEE 2030.5 defines many optional components that aren't required to meet Australian use-cases CSIP-AUS spells out which parts of the protocol are needed, and where there are multiple ways of doing something, provides guidance or restrictions on those options.

To this end, CSIP-AUS:

- Specifies the parts of IEEE 2030.5 that apply in complying Australian implementations and the ways in which they should be implemented.
- Defines several new extensions that enable Australian requirements not supported by IEEE 2030.5, most notably components necessary to communicate DOE limits.

CSIP-AUS also has a relationship with CSIP (the *Common Smart Inverter Profile*), which is a separate implementation guide for IEEE 2030.5. CSIP was originally developed by a collaboration of California-based utilities alongside local industry participants, and uses the IEEE 2030.5 protocol to enable utility control of customer DER. As CSIP-AUS builds on the extensive detail within CSIP, the two documents need to be read in conjunction to attain a full understanding of Australia requirements. Areas where CSIP-AUS differs from CSIP have bolded text, as shown below:

Table 1: Comparison of CSIP and CSIP-AUS content

<u>CSIP</u> (lines 177 – 183)	<u>CSIP-AUS</u> (lines 177-183)
Note that the notion of a DER in CSIP is a logical concept generally thought of as one or more physical inverters organized and operating as a single system with a common point of aggregation behind a single point of common coupling (PCC) with the utility. This allows the management of a plant/system possessing a single PCC regardless of whether it is composed of a single inverter or many. It is the responsibility of the aggregator system to manage the underlying inverters to meet the requirements of the settings provided by the utility server. The specific interpretation of the DER being a single entity or a related group is established at the time of interconnection with the utility.	Note that the notion of a DER in CSIP is a logical concept generally thought of as one or more physical inverters organized and operating as a single system with a common point of aggregation behind a single point of common coupling (PCC) with the utility. This allows the management of a plant/system possessing a single PCC regardless of whether it is composed of a single inverter or many. It is the responsibility of the management system to manage the underlying inverters to meet the requirements of the settings provided by the utility server. **A DER client at each site must be able to control and measure the site power export at the PCC. DERs on a site behind the point of common coupling may be controlled independently, however only one DER on a site SHALL be used to control and measure the import and export power to the grid.**

2.4 Potential future extensions to CSIP-AUS

The DERIAPTWG is continuing to develop CSIP-AUS to incorporate new features and functionality, improve technical aspects of the profile, and align with developments in IEEE 2030.5, CSIP and industry expectations. With the publication of CSIP-AUS version 1.2 expected imminently focus is now on development of version 1.3, which is expected to incorporate the following new capabilities –

- Development of an optional Pricing module which will support the communication of both DNSP and retailer pricing/tariff information dynamically over the protocol.

- Incorporating broader aspects of IEEE 2030.5 that will enable the optional dispatch of VPP devices such as aggregated energy storage assets.
- Ability to read other DER settings such as volt-var to confirm Australia A settings have been applied
- A wide range of technical improvements aimed at enhancing scalability, robustness and consistency of implementations.

Beyond version 1.3 there are long-term plans to fully redraft CSIP-AUS – both to fully revise technical requirements to take consideration of the many learnings and improvements that have arisen since the profile was first conceived, and to incorporate the numerous improvements to the base IEEE 2030.5 protocol included in the 2023 of that standard. This revision is tentatively being referred to as *CSIP-AUS 2.0* and may coincide with elevation of SAHB 218 to full Australian Standard (subject to Standards Australia's endorsement of this proposal).

2.5 How CSIP-AUS is being used

This list provides some indications of jurisdictions using CSIP-AUS at time of publication. This is not an exhaustive and may not be accurate in the future.

State	Usage
South Australia	Used for dynamic export limits[2] and can be used for the relevant agent framework[3].
New South Wales	Used in dynamic export limit trials and will be used in the upcoming solar emergency backstop in 2026[4].
Victoria	Used in dynamic export limit trials and is being used in the solar emergency backstop[5].
Queensland	Offered for dynamic import and export limits[6].
Western Australia	Used in emergency solar management[7] and for VPP product participation[8].

[2] https://www.energymining.sa.gov.au/industry/hydrogen-and-renewable-energy/solar-batteries-and-smarter-homes/regulatory-changes-for-smarter-homes/dynamic-export-limits-requirement

[3] https://www.energymining.sa.gov.au/industry/hydrogen-and-renewable-energy/solar-batteries-and-smarter-homes/regulatory-changes-for-smarter-homes/information-for-relevant-agents

[4] https://www.energy.nsw.gov.au/nsw-plans-and-progress/regulation-and-policy/public-consultations/solar-emergency-backstop

[5] https://www.energy.vic.gov.au/households/victorias-emergency-backstop-mechanism-for-solar

[6] https://www.energex.com.au/our-services/connections/residential-and-commercial-connections/solar-connections-and-other-technologies/dynamic-connections-for-energy-exports/about-dynamic-connections

[7] https://www.wa.gov.au/organisation/energy-policy-wa/emergency-solar-management

[8] https://www.synergy.net.au/Our-energy/Household-energy-assets/DER-Industry-Resources

ACT	Used in dynamic export limit trials[9] and proposed being used in the upcoming solar emergency backstop[10].

3 CSIP-AUS USE CASES

CSIP-AUS has been designed and developed to address a core set of use cases that relate to onboarding and remote control and management of CER to manage network, retailer and system impacts. The key uses cases are outlined below.

3.1 Emergency Backstop

Emergency backstop requirements are being imposed by Australian governments to ensure the power system can remain stable during periods of very high solar generation. The “mild spring day” scenario, when there is plentiful solar generation with little heating or cooling demand, is the typical example of where this use case might be required.

The requirements for this emergency measure are set by the Australian Energy Market Operator (AEMO). AEMO determines Minimum System Load (MSL) requirements for each network area and communicates these to DNSPs for implementation as they are forecast to occur.

CSIP-AUS provides a range of tools to assist DNSPs to meet their emergency backstop obligations for managing MSL. Each of the following management signals is supported, and can be utilised in isolation or in combination by DNSPs, depending on the specific network and jurisdictional requirements:

- Zero export limit – DER receiving this signal must reduce generation to ensure no active power is exported from the site.
- Zero generation limit – DER receiving this signal must reduce generation to ensure there is no net generation from DER within the site.
- De-energise – DER receiving this signal must cease to generate and consume power.
- Disconnect – DER receiving this signal must disconnect from their AC supply.

The exact requirements for emergency backstop vary from one jurisdiction (e.g. one state) to another. This includes both the way the use case is implemented and the inclusions and exclusions that apply.

Scenarios can be complex, and controls may not cover all generation devices. For example, ‘de-energise’ and ‘disconnect’ operational modes may not be appropriate for batteries (even though these are generation devices) as it would prevent these devices from consuming power from the grid at the critical time.

Grandfathering arrangements also vary from one jurisdiction to another, as authorities take a variety of approaches to allow customers to make best use of their investment in previous installations.

[9] https://arena.gov.au/projects/project-converge-act-distributed-energy-resources-demonstration-pilot/

[10] https://hdp-au-prod-app-act-yoursay-files.s3.ap-southeast-2.amazonaws.com/5317/4216/9923/2025_Emergency_Backstop_Capability_Consultation_Paper_FA_access.pdf

Controls such as de-energise and disconnect do not generally allow self-consumption for the generating DER under control. They are considered a last resort after reductions from other means (such as export limits) have been exhausted.

3.2 Flexible Exports

Even in cases where there is no emergency condition, it may be useful for a DNSP to place limits on the amount of power that can be exported from a site. In some cases, this may prevent the onset of an emergency condition. It may also be the case that an energy retailer may want to limit or allow different levels of power export in order to respond to the demands of the National Energy Market.

In these circumstances, CSIP-AUS provides controls that will allow DNSPs and/or retailers to control export from a customer's site. Controls include:

- Maximum export limit – DER receiving this signal, delivered as a DOE control, must limit export of power to within the limits specified in the envelope. This limit could be zero, as for the Emergency Backstop case, but could be a higher number, depending on network conditions.
- Maximum generation limit – DER receiving this signal, also delivered as a DOE control, must limit net generation at the site to within the bounds set in the DOE.

A common way to implement this use case is to allow flexible exports to be generally higher than static limits. For static export limits, the network must be conditioned to allow power up to this limit to be exported at any time. With flexible export limits, the network can allow much more power to be exported or generated at certain times, with the understanding that these limits can be much stricter in times when exported power is undesirable.

3.3 Flexible Demand Reduction

An emerging use-case is to extend CSIP-AUS-based CER management schemes to flexible load devices to mitigate the need for rotational load shedding[11] when there is insufficient generation at the bulk power system level to meet demand (i.e., under lack of reserve level 3 (LOR3) conditions). Whereas DOEs (generated and issued by DNSPs) are designed to manage localised network constraints, emergency backstop capabilities are designed to manage system security risks identified by AEMO.

3.4 Microgrid Controls

A microgrid (MG) is an embedded electrical network of interconnected loads and distributed energy resources (DER) with the capability to operate grid-connected or islanded from the grid. It can support bi-directional power flows through the point of common coupling (PCC).

Typically, operation of a microgrid may include some or all the following objectives:

- Least cost dispatch (economic dispatch), which minimises grid energy or fuel consumption and maximises renewable energy production
- Store energy during daytime (i.e., from rooftop PV systems) to use at night or times with no sunlight
- Regulate (control) the system voltage and frequency to maintain power quality standards
- Manage the peak demand due to e.g., HVAC, using controlled loads
- Mitigate grid instability due to renewable energy output intermittency

[11] Rotational load shedding temporarily deenergises areas of a distribution network on a rotating basis (i.e. a localised 'blackout') so as to meet overall load reduction targets set by AEMO.

- Seamless transition to islanded operation if the grid fails, and automatic re-synchronisation when the grid returns

The use of CSIP-AUS to control microgrids can be cheaper and more flexible than traditional SCADA approaches.

CSIP-AUS can be used at the PCC to monitor the state of the microgrid as a whole and its impact on the connected grid.

CSIP-AUS can also be used to effect controls for the microgrid. These will typically be DOE-based to control the import and export of power from the microgrid but are likely to also include generation and consumption level controls that do not support self-consumption. These are similar to those used in previous use cases, but several controls might be used in combination.

3.5 Additional Use Cases

It is important to note that the previously outlined use cases intend to work alongside other use cases such as Virtual Power Plants (VPPs) or Price-Responsive Resources and does not preclude their use at the same site. Where signals between use cases are conflicting, the requirements of the connection contract set by the DNSP is considered higher priority.

The "Integrating price-responsive resources into the NEM" rule change establish requirements for DNSPs to consult with affected Voluntarily Scheduled Resource Providers, regarding how those limitations could be incorporated into dispatch bids.

4 COMMON INTERACTIONS

The CSIP-AUS (HB 218) covers the following common interactions for clients:

- Discovery – the process for a new EndDevice to be connected to a utility server for the first time and set up for regular operations.
- Regular operations – interactions that are performed between the server and client on an ongoing basis.
- On change – actions that are performed by a client when a relevant change is detected/enacted.

The interactions between the server and the client are presented in Figure 2 and Figure 3 below:

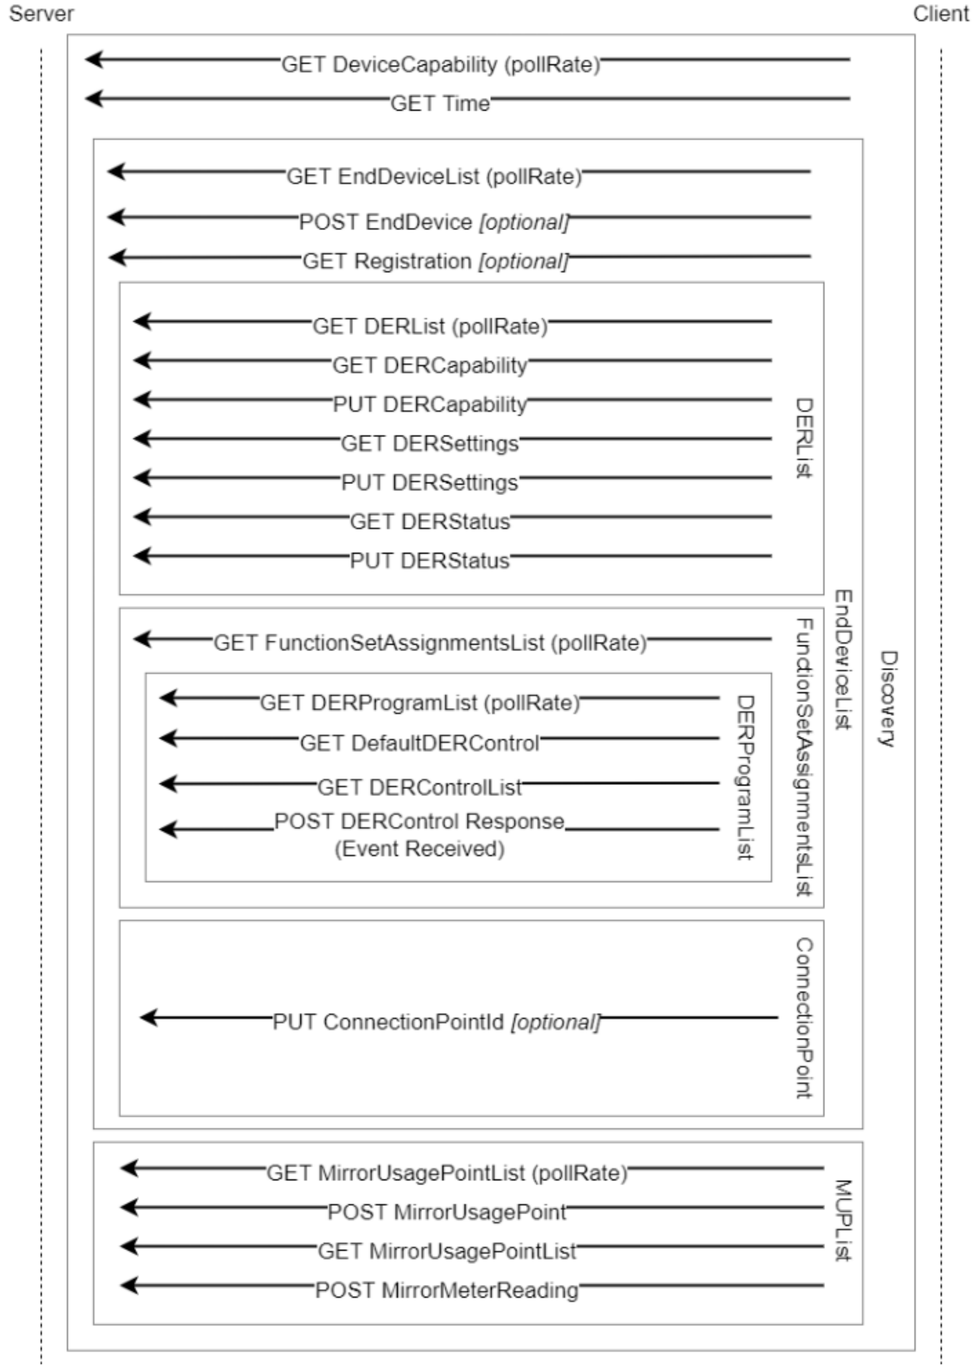

Figure 2: CSIP-AUS communications interactions – Discovery

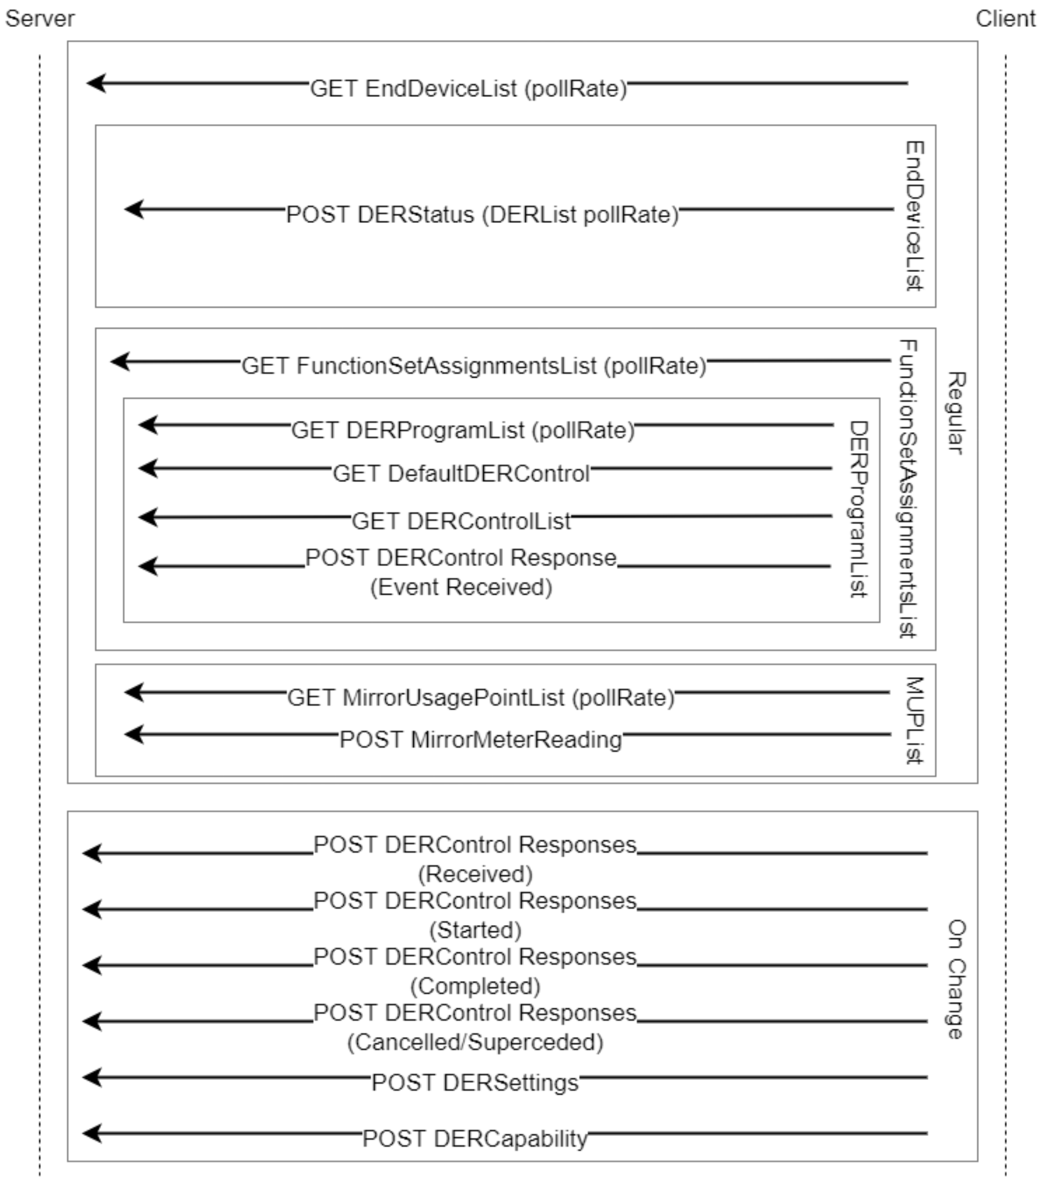

Figure 3: CSIP-AUS communications interactions – ongoing

4.1 Onboarding/Setup

4.1.1 Connection to a utility server

Details for connecting to a utility sever will be provided through an 'out of band' process by each utility server operator (either in their handbook or as part of the onboarding process). For each utility server instance (e.g. sandbox, certification, production) this will include the location of the device capability resource and the protocol and port required to connect.

As clients may be operating with more than one utility server in Australia, clients should retain an understanding of each utility server that they can connect to and select the right utility server at the time of installation. When connecting to a network operator's utility server, it is recommended that clients use the NMI to identify who the network operator is, and map this to the location of the server.[12]

[12] AEMO NMI Allocation List

4.1.2 Certificate Provisioning

Certificates shall be provisioned through a Certificate Signing Request workflow. Examples of this workflow are provided in Section 0.

4.1.3 HTTP Response Codes

Specific HTTP response codes (particularly 4XX response codes) are often used by utility servers to inform a client of the source of a particular issue. The codes listed in this section are used by some utility servers to flag specific circumstances. It should be recognised that not all of these are supported by all servers, and so these should be treated as informational rather than relied upon in all circumstances. Future updates to CSIP-AUS may make these requirements mandatory.

Technical detail on the general meaning of HTTP responses in general web use can be found in section 15.5 of RFC 9110[13]. The Mozilla Foundation also provides helpful guidance[14] on the practical meaning of a range of HTTP response codes in a web context.

Some of the common responses are included below, but it is not considered an exhaustive list. Further examples of how these and other HTTP response codes should be interpreted can be found in Section **Error! Reference source not found. Error! Reference source not found.**.

4.1.3.1 Successful Responses

A response in the 200-299 range represents a successful response.

4.1.3.2 Client Error Responses

A response in the 400-499 range represents an unsuccessful response.

- 400 – This specific resource creation has failed for some reason, stop trying with this resource details or payload.
- 401 – Unauthenticated/Unauthorised. This could be a result of trying to access the wrong resource, or indicate an issue with the client certificate; however, it is more likely the TCP connection will be dropped completely rather then receiving this response.
- 403 – This client is not currently authorised to create any new instances of this resource. Please contact the Utility Server operator.
- 404 – The requested resource does not exist. The utility server may return this response for:
    - A valid URL if the client does not have access, but it does not want to reveal the existence of the resource.
    - Where a link to a resource is out of date, and the resource has been moved to a new location. In this situation clients should attempt to re-retrieve the end-point URI from the server by accessing another resource (e.g. a parent resource or list) that contains the relevant link.
    - Where a resource that is expected to exist, has not yet been created by the server. In this case the client should retry periodically to retrieve to resource with exponential backoff.
- 409 – This resource is a duplicate. The utility server may alternatively return a 400 in this scenario if it does not want to reveal the existence of the resource.

[13] https://datatracker.ietf.org/doc/html/rfc9110#name-status-codes

[14] https://developer.mozilla.org/en-US/docs/Web/HTTP/Status

- 429 – This indicates that you are being rate limited by the server as it has received too many requests (may be a global request limit, not necessarily caused by your client) and you should retry the request after a suitable delay.

4.1.3.3 Server Error Responses

A response in the 500-599 range represents an issue with the server. There is nothing the client can do to resolve these issues, however the client should retry the request after a suitable delay for transient issues.

For example, a 503 response may indicate that the server is down for maintenance and may include a Retry-After HTTP header with an estimated time of recovery.

4.1.4 Registration

4.1.4.1 Out of Band Registration

Utility Clients must support out-of-band registration (server creates the EndDevice and all subordinate resources).

Utility Servers must support out-of-band registration for all EndDevices. For aggregator clients specifically, prior to connecting to the Utility Server for the first time, the Utility Server should create an EndDevice to represent the aggregator and configure that EndDevice with suitable permissions. When aggregators retrieve their EndDeviceList for the first time it will include a single EndDevice that represents the Aggregator Client. Once connected aggregator clients will need to have EndDevices registered (in-band or out-of-band) for each EndDevice that they are managing.

4.1.4.2 Inband Registration

Utility Clients should support in-band registration of EndDevices and ConnectionPoints through the process described in SA HB 218 section 4.2 and section 11. For direct clients this registration of the device and association with the ConnectionPoint is only required once, for the EndDevice that the client is representing and the associated ConnectionPoint.

Utility Servers must support in-band registration for all EndDevices that are not the EndDevice representing the aggregator. Utility Servers may choose to refuse an in-band registration request for an EndDevice for any reason and will provide a 4XX response to the client. The client should understand the response codes as follows and take appropriate corrective action as specified in section 4.1.3.

In the future it is expected that in-band registration will be used for all almost all EndDevice registrations. There may be specific cases where out-of-band registration is required e.g. onboarding of proxy/aggregator “virtual EndDevice” described in CSIP section 6.1.5.

4.1.4.3 Connection Point Registration

ConnectionPoint registrations should be performed immediately following EndDevice registrations as Utility Servers may archive EndDevices that are not associated to a ConnectionPoint after a certain duration.

Utility Servers may choose to refuse an in-band registration request for a Connection Point for any reason and will provide a 4XX response to the client. The client should understand the response codes as follows and take appropriate corrective action. Specific response codes relevant for Connection Point registration are listed below, and detailed in Section 4.1.3.

- 400
- 403

- 409

4.1.5 Resource Instantiation

As part of the EndDevice instantiation, Utility Servers will create links and instances for all relevant resources. This may include the following resources:

- EndDevice
- FunctionSetAssignmentList
- FunctionSetAssignment(s)
- SubscriptionList
- Registration
- DERProgramList
- DERProgram
- DERControlList
- DERControl
- DefaultDERControl
- DERList
- DER[15]
- DERStatus
- DERSettings
- DERCapability

Clients are not required to create any of the resources in the list above.

4.1.6 DERCapability and DERSettings

As part EndDevice onboarding Clients should make submissions to DERCapability and DERSettings. Clients should submit all required values for DERCapability and DERSettings in a single request to each resource.

4.1.7 MirrorUsagePoints and MirrorMeterReadings Management

As a part of the onboarding workflow clients shall create relevant MirrorUsagePoints (MUPs) and MirrorMeterReadings (MMRs) as defined in SA HB 218 Section 8 for each EndDevice under management.

When creating MUPs and MMRs, clients solutions shall implement measures to avoid collisions with existing mRIDs under the same organisation. Where a collision is detected, clients shall re-generate a new mRID to ensure that the MUPs and MMRs can be created.

Generated mRIDs need to have an even number of characters, and in most cases must be 32 characters. The PEN should be included in generated mRIDs to ensure uniqueness:

e.g. 0xFFFFFFFFFFFFFFFFFFFFFFFF[XXXXXXXX], where [XXXXXXXX] is the PEN

The characters should also be in uppercase and only include the characters A-F and 0-9.

Clients should create a single MUP for site level measurements and a single MUP for device level measurements. For each MUP created, clients can create multiple MMRs for each of the relevant monitoring points to be reported.

---

[15] In the CSIP-AUS implementation, there shall be a single DER per EndDevice.

Clients should create a single MMR for each readingType that is being submitted to the server and continue to use that MMR while it remains active.

Clients may create a single MirrorUsagePoint for each telemetry stream definition, however this approach is not preferred as it utilises additional resources unnecessarily.

When reconnecting to the server under any loss of connection, clients should retrieve all MUPs and MMRs and confirm the presence/absence of the relevant MUP/MMR for each EndDevice under management.

Clients should only create new MUPs and/or MMRs, where a required MUP or MMR has not been created or has been removed rather than recreating an existing MirrorUsagePoint as a matter of course. For example, a server may timeout if it does not receive telemetry from a client after 72 hours, and delete the respective MirrorUsagePoint or remove it from the MirrorUsagePointList.

Utility Servers may choose to refuse an in-band registration request for a Connection Point for any reason and will provide a 4XX response to the client. The client should understand the response codes as follows and take appropriate corrective action. Specific response codes relevant for Connection Point registration are listed below, and detailed in Section 4.1.3.

4.1.8 Subscriptions

Where a 2030.5 resource is subscribable, aggregator clients should implement subscription for that resource to reduce traffic on both sides.

Subscriptions should trigger when any object on that resource have changed. For example, if an event is cancelled that event should be resent with the updated status. Note that for performance reasons it is better if the server only resends events that have had an update and not the entire list (i.e. don’t resend unchanged events). Clients should treat any received updates as a delta from previous polls or subscriptions.

4.1.8.1 Management of Subscriptions

Clients and servers should be aware of the recommendations in IEEE 2030.5 relating to good practice for maintenance of subscriptions, as detailed in IEEE 2030.5 Section 8.7.3. Utility interconnection handbooks may also apply additional recommendations or requirements that client implementers should be aware of – for example some utility servers expect clients to resubscribe at least every 24 hours (and no faster than hourly).

4.2 Ongoing Operations

4.2.1 Post and poll rates

The post rates and poll rates for resources are generally specified by the server, and it is the responsibility of the client to implement the specified post rates and poll rates. While some resource have post rates and poll rates specified against themselves, other resources inherit post/poll rates from parent resources. The table below maps the poll rate/post rate that a given resource should use, particularly when these rates are defined on a separate resource. If a rate is not provided by the utility server the default rate in IEEE 2030.5 of 900 seconds will be used. Often slower poll rates (e.g. 24 hours) are used for higher-level resources that are not expected to change frequently, to reduce unnecessary load on servers and clients.

Implement resource	Post or poll rate	Resource the rate is defined on
DeviceCapability	Poll	DeviceCapability pollRate
EndDeviceList	Poll	EndDeviceList pollRate
EndDevice	Poll	EndDeviceList pollRate
Registration	Poll	Registration pollRate
MirrorUsagePointList	Poll	MirrorUsagePointList pollRate
FunctionSetAssignmentsList	Poll	FunctionSetAssignmentsList pollRate
SubscriptionList	Poll	SubscriptionList pollRate
Time	Poll	Time pollRate
DERProgramList	Poll	DERProgramList pollRate
DERProgram	Poll	DERProgramList pollRate
DERControlList	Poll	DERProgramList pollRate
DERControl	Poll	DERProgramList pollRate
DefaultDERControl	Poll	DERProgramList pollRate
MirrorUsagePoint	Poll	MirrorUsagePointList pollRate
DERList	Poll	DERList pollRate
DERStatus	Post	DERList pollRate and on change
ConnectionPoint	Post	On setup
DERSettings	Post	On change
DERCapability	Post	On change
DERControlResponse	Post	On change
MirrorUsagePoint	Post	MirrorUsagePoint postRate

4.2.2 Retrieving End-point Links

Clients should refresh from server/resource links and update internal state where applicable on a regular basis, as determined by the poll rate of the relevant resource. A prime example is the DeviceCapability resource containing links for the EndDeviceList and MirrorUsagePointList end-points.

4.2.3 Retrieving Active DERControls

Clients are expected to access FunctionSetAssignments, DERPrograms and their subordinate default and active DERControls at the respective polling rates. For FunctionSetAssignments this is the polling rate defined on the FunctionSetAssignmentsList resource. eFor DERPrograms and DERControls this is the polling rate defined on the DERProgram resource. Where polling rates are not specified the default rate of 5 minutes should be used as described in CSIP-AUS / SA HB 218 section 4.5.

4.2.4 Reporting Telemetry

When submitting readings, the measurement aggregation period shall be equal to the MirrorUsagePoint post rate specified by the server and aligned to the UTC tick for the relevant post rate. For example, for a five-minute post rate, measurements shall be aligned to the five-minute tick, for a one-minute post rate, measurements shall be aligned to the one-minute tick.

For measurement and calculation of average real and reactive power, the preferred approach is for the hardware to measure the total energy produced/consumed over the period, and divide that value by the period duration. If this approach cannot be implemented a permissible alternative is to have the values calculated from sampled values across the measurement period. The sampling frequency should be once per 200 ms but shall be at least once per 1 s. For each sample, the value shall be measured using the approach in AS/NZS 4777.2 section 2.13.

Any measurement taken either at the DER or at the site level shall meet the accuracy requirements specified in section 2 of AS/NZS 4777.2 and documented in section 7 of AS/NZS 4777.2.

Signing conventions are used by default for providing telemetry and align to the sign convention specifications in IEEE 2030.5. For telemetry at the DER level, active sign convention applies, i.e. generation is positive and load is negative, For telemetry at the site level passive sign convention applies i.e. export to the grid is negative and import from the grid is positive.

Telemetry reading values are a signed Int16 (-32768 to +32767). Values larger than this should make use of the PowerOfTenMultiplier attribute.

4.3 On change

4.3.1 Reporting DERCapability and DERSettings

When the physical values of asset ratings or adjusted ratings change, the client shall make a PUT request to DERCapability and/or DERSettings to reflect the changes in those capabilities/settings. All values for DERCapability and DERSettings shall be included in the same request.

Definition for the attributes under DERCapabilities and DERSettings can be found in Section 10.10.4.4 of IEEE 2030.5. In some cases CSIP-AUS expects a slightly different interpretation for specific attributes.

For EndDevices managing multiple independent CER, the aggregation method for specific values of both DERCapability and DERSettings are as follows:

- modesSupported and modesEnabled shall reflect the logical OR of all managed DER;
- Active, reactive and apparent power and current ratings & settings shall reflect the sum of the individual ratings;
- Type shall reflect the mix of DER under management e.g. Type 1 “Virtual of mixed DER” for complex mixes of a range of DER types.

- Voltage ratings shall reflect the minimum of rtgMaxV/setMaxV and maximum of rtgMinV/setMinV of the individual ratings;
- updatedTime shall reflect the time of the most-recent setting change.

4.3.2 Reporting DERStatus

Clients shall make submissions to the DERStatus resource for each EndDevice that they are representing/managing. Further guidance on the physical observations that should elicit certain status responses are provided for OperationalModeStatus and genConnectStatus.

operationalModeStatus:OperationalModeStatusType

1 - Off: This status shall be reported whenever the conditions for “Operational mode” are not met.

2 - Operational mode: This status shall be reported where one or more components of the EndDevice are energised from a DC energy source.

genConnectStatus:ConnectStatusType

Bit 0 – Connected: This bit shall be asserted where at least one DER with a generation source that is part of the EndDevice has a live AC connection to the grid. This does not necessarily mean that the EndDevice is operating or available.

Bit 1 – Available: This bit shall be asserted where at least one DER with a generation source that is part of the EndDevice is online, providing telemetry and capable of receiving and responding to controls.

Bit 2 – Operating: This bit shall be asserted where there are active power flows on either AC-DC or DC-DC circuits on at least one DER with a generation source that is part of the EndDevice.

Bit 4 – Fault/Error: This bit shall be asserted where any internal fault/error is detected that prevents any DER with a generation source that is part of the EndDevice from responding or being capable of responding to controls, inclusive of communications errors between the asset and Utility Client.

4.3.3 Submitting DERControlResponse

Clients shall submit DERControlResponses at the location specified in the replyTo attribute of the DERControlEvent. Responses shall be submitted by the client on change when the criteria for each response category are met, and shall be submitted to the server within five minutes of the criteria being met.

1- Event Received. The event has been received by the client. Note: As clients are not necessarily hosted locally, an event will be marked as received even if there are no active communications to the site.
2- Event Started. The DER has started responding to the event.
3- Event Completed. The DER has successfully completed the event.
6- Event Cancelled. The cancellation request has been received by the client. Note: As clients are not necessarily hosted locally, an event will be marked as received even if there are no active communications to the site.
7- Event Superseded. The client has received another event that supersedes the existing event. Note: As clients are not necessarily hosted locally, an event will be marked as superseded even if there are no active communications to the site.

4.3.4 EndDevice Removal

Utility Clients will not have the ability to remove/delete EndDevices through an in-band workflow. The utility server operator will provide an out of band mechanism for EndDevices to be removed/deleted.

5 ABNORMAL OPERATIONS

5.1 Edge Cases

5.1.1 Server Loss of EndDevice Resources

In some circumstances a utility server may lose internal references to resources associated with one or more EndDevices (including resources subordinate to the EndDevice resource). In extreme cases – for example if a server's database is rebuilt – a utility server may lose references to *all* resources obtained via in-band mechanisms. In these situations list resources that should include information related to the client – for example the EndDeviceList – would not include the respective client information. Related issues may also result in a previously operational HTTP end-point returning 404 errors rather than a valid response.

Clients should be capable of detecting and adapting to these situations. In general if a client detects that a list resource no longer contains expected information that the client previously provided, the client should recreate that resource via the expected means e.g. by POSTing to the EndDeviceListLink with the appropriate information. This may also require the client to update mRIDs for the new resources.

Clients receiving a 404 error should attempt to retrieve the correct URI by accessing the resource that would contain the respective link as described in section **Error! Reference source not found.** above.

5.1.2 Inactive EndDevices

EndDevices can become inactive for a variety of reasons, including decommissioning of a device, a DER changing its CSIP-AUS client, or a site with DER changing ownership. This can result in circumstances where data operations continue relating to an EndDevice that in effect no longer exists, resulting in wasted HTTP exchanges and unnecessary costs.

In general it is the utility server's responsibility for resolving situations like these when they occur. This may include the utility server sending an appropriate signal to disable the client from representing the relevant EndDevice. It is also considered best-practice for clients and other client-side actors to endeavour to make utility servers aware of these situations so they can take appropriate action.

5.1.2.1 EndDevice LifeCycle

A utility server may implement practice where EndDevices are removed periodically as a matter of standard operation to ensure that an EndDevice doesn't remain live after the DER it represents has been decommissioned. Under this model a client would be expected to detect the absence of the EndDevice from the EndDeviceList provided by the server, and recreate it as described in section 5.1.1 above.

5.2 Loss of Communications

Specific behaviours for each of these circumstances are included below.

5.2.1 Loss of communications between client and utility server

Loss of communications includes outages (or unexpected error messages) on the server side or outages on the client side. Clients should continue to attempt to connect to the server when the client's physical communication becomes available, and make requests against the relevant server resources at the specified poll rates for normal operations.

In future schedules should taper towards default control (particularly to e.g. overnight timeframes) to avoid communications outages leading to e.g. step-changes in generation when fleets of DER drop to default control values.

If a widespread internet outage coincides with a change in forecast network conditions (for example clouds reducing solar generation) the network must be operated safely even though it cannot send new controls to the offline clients. Due to this uncertainty surrounding forecasts, utilities should taper or decay forecasts such that they trend towards the default control at the end of the forecast period. For example

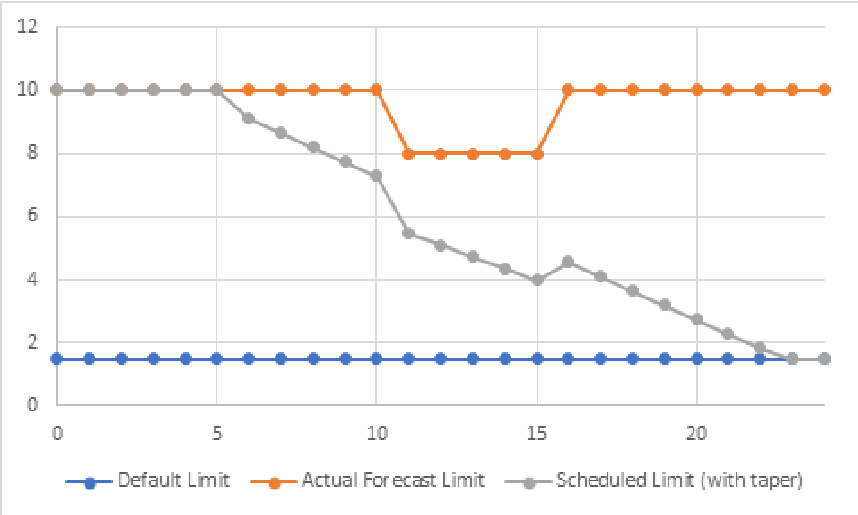

Utility servers will utilise the DefaultDERControl to specify the desired failsafe behaviour of the DER client should a communications failure occur and after all DERControl events have been exhausted.

It is a CSIP requirement that clients store at least 24 scheduled events for each DER. Should there be a loss of communications, sites must complete any locally scheduled events until exhausted, and then revert to default settings (or fixed limits if a default is not specified).

If capable of doing so, clients should continue to buffer telemetry for the duration of the outage.

After the client has reestablished communications with the utility server, it should attempt to retrieve the most up-to-date schedule of DER controls and enact them. It should also provide updated operating data including status, alarm and telemetry information (including any telemetry that was buffered during the outage).

5.2.2 Loss of communications between client and managed DER

DER that lose communications with their CSIP-AUS communications client are expected to still be able to meet the minimum operational requirements of the profile while communications are down. This can be achieved by ramping down to default control values and remaining there until the link with the client has been restored.

Under no circumstances should DER immediately revert to default controls (without following ramping limits or not waiting until the end of the active control) as this could result in large power swings on the network for customers with unreliable internet connections.

5.3 Network Outages

5.3.1 Re-establishment of communications following power outage

Following a disconnection of the DER from AC supply (which may/may not have been caused by a power outage), the client MUST remove all existing *DERControls*, revert to the *DefaultDERControl* and re-poll all associated *DERPrograms* to establish if any *DERControl* events have been added or changed whilst offline.

This process shall be applied as a secondary process in addition to the connection and reconnection procedure defined in AS/NZS 4777.2, which includes waiting for 1 minute after power is restored before generating as well as ramping limits.

5.3.2 Site running islanded

When a site is running islanded (*alternative supply* in AS/NZS 4777.2) and not connected to the grid, it is not expected that systems will follow controls sent over CSIP-AUS. The requirements in section 5.3.1 apply when the site re-connects to the grid.

6 DER RESPONSE

6.1 Ramp rates

The ramp rate for transitioning to an active control is specified by the rampTms attribute of the upcoming control. If absent, the default ramp rate (setGradW) applies for transitioning to the active control.

When transitioning from an active control back to a default control, the setGradW specified in the default control applies.

Within AS4777.2:2020 there are two ramp rate references applicable when interpreting CSIP-AUS response. These are:

- Section 3.3.4.2 Gradient of power rate limit. 16.67 % of rated power per minute with a nominal ramp rate time of 6 minutes. Within CSIP-AUS this is referred to as the default ramp rate.
- Section 6. Soft limit generation and export limit control that an inverter or multiple inverter combinations shall monitor and control the output within 15s.

The Gradient power rate limit applies when:

- There is no setGradW or rampTms AND
    - Moving to a default control from an active control at the end of an active control and no further scheduled control
    - On a communications outage and reverting to default control

The soft limit applies when:

- There is no rampTms
- When controlling to an active control
- Whilst managing to a current default control. E.g. Output of an inverter or multiple inverters response whilst no active controls and default export control of 0, site must respond within 15s.

If setGradW has not been specified in the default control, it should be assigned as per the ramp rate specified in AS 4777 (0.278% per second). For systems with multiple DER, the ramp rate applies to the total aggregate DER capacity when converting to kW, as per the definition of setMaxW.

For systems with multiple DER, the ramp rate applies to the total aggregate DER response when converting to kW.

For implementation, it is likely that the setGradW will need to be converted to W/s. For example, a site with a setMaxW of 10,000 W and a setGradW of 28, would have a default ramp rate of 28 W/s (or 1.7 kW/min). A transition to an active limit of 10 kW from 1.5 kW with a rampTms of 60s would have a ramp rate of 142 W/s (8,500/60).

In an example scenario shown in

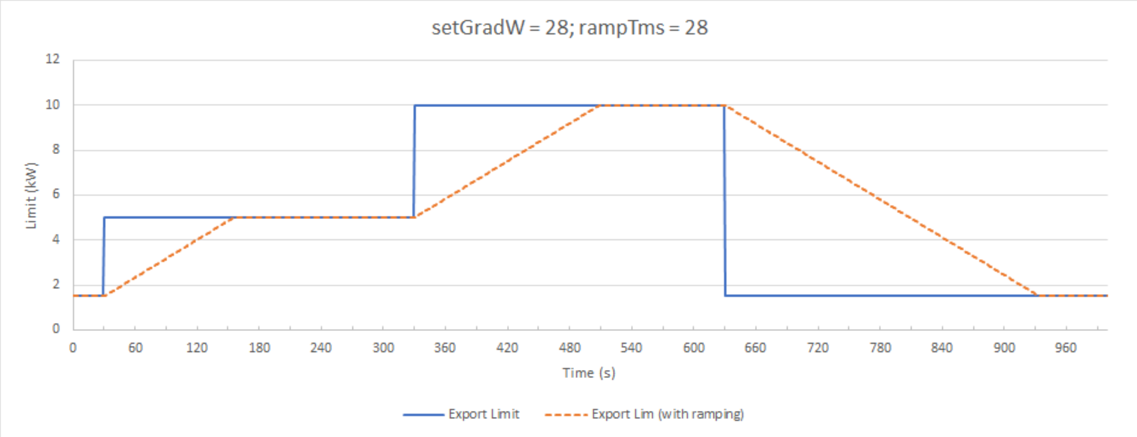

Figure 4 and Figure 5, there is a 10 kW DER, with a default export limit of 1.5 kW. There is an active control of 5 kW from 30-350s, and an active control of 10 kW from 350-650s. At 650s the system will revert to the default limit as there are no further events.

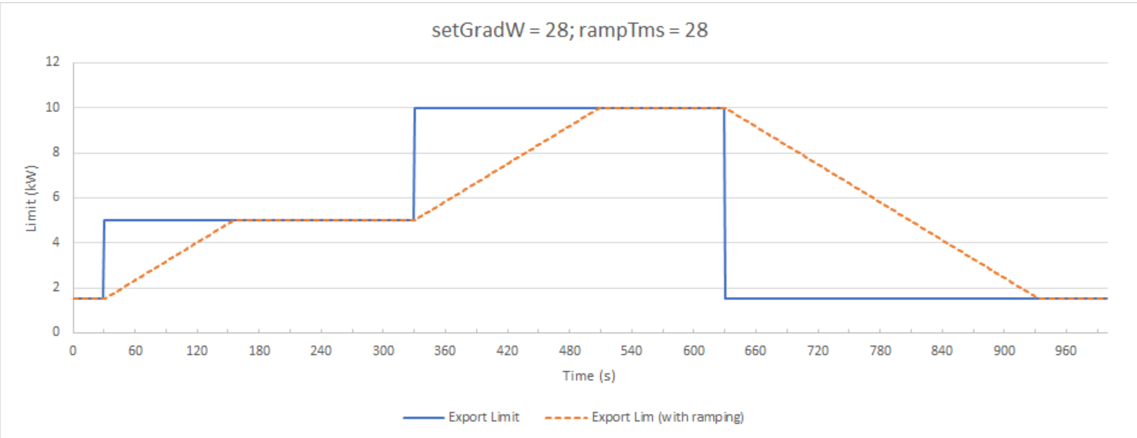

Figure 4: Example where rampTms and setGradW are the same.

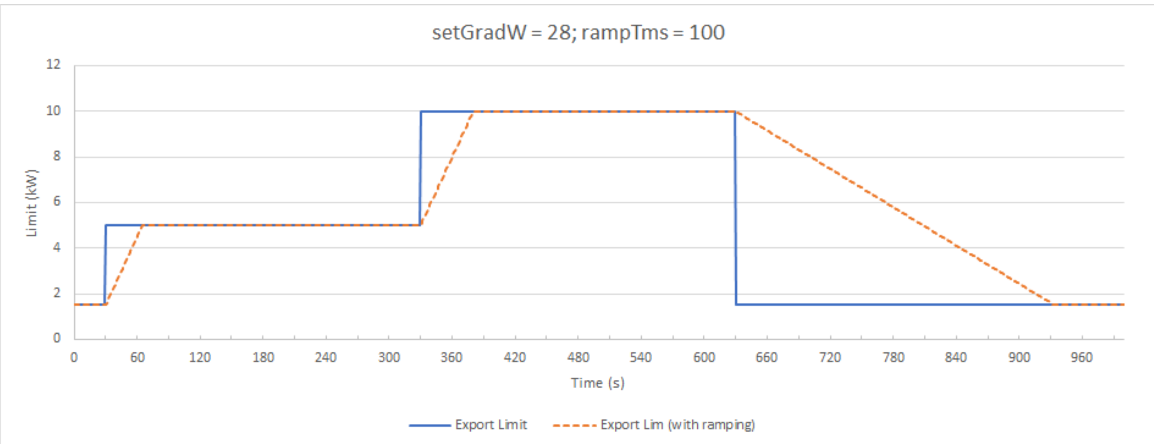

Figure 5: Example where rampTms and setGradW differ.

It is noted that ramp rates must be complied with whenever possible but may not be possible for all DER types (such as DER controlled only by an on/off switch). As the values are limits – all values less than the ramped limit are deemed compliant.

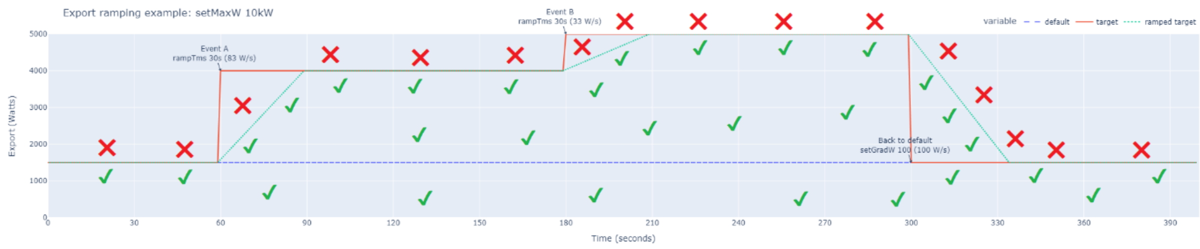

Figure 6: Example showing how limits can be less than the limit

6.2 Active limits

DER shall be able to adhere to active limits that are provided to their client by the utility server.

In general clients and DER are required to enact any controls they receive that claim to support and have enabled via modesEnabled and doeModesEnabled.

At a minimum all CSIP-AUS DER are expected to support the opModConnect and opModEnergize controls, with support for other CSIP-AUS controls depending on the type of DER. Generators would normally be required to meet the opModExpLimW (export limit) and opModGenLimW (generation limit) controls, while load devices would normally be required to meet the opModImpLimW (import limit) and opModLoadLimW (load limit) controls. Energy storage can choose either or both of export/generation or import/load limits to meet, depending on any requirements that may be applied to them based on the jurisdiction or network they operate in.

All CSIP-AUS limits apply to active power at a site level, as shown in Figure 7. This means:

- opModExpLimW effects a limit on the total active power that can be exported from a site. This may mean that the generation on-site can vary in response to local load fluctuations, provided that the export limit is not exceeded.
- opModGenLimW effects a limit on the total active power that can be generated on a site by controllable DER.
    - This limit applies as a net sum to the aggregation of all DER on a site that are being managed as part of the EndDevice, *including controlled loads*. Specifically this means that if a managed load device (for example a battery or EV charger) is consuming active power, the remaining DER on site can generate additional power to offset this while still meeting the generation limit. Figure 7 illustrates this by showing the point at which opModGenLimW applies.
- opModImpLimW effects a limit on the total active power that can be imported by a site. This may mean that the load on-site can vary in response to local generation fluctuations, provided that the import limit is not exceeded.
    - Where an EndDevice has been issued an import limit control, only controllable DER that are part of the EndDevice are expected to change their behaviour to respond to the limit, reducing load to zero where the import is exceeded. Other loads on site that are not managed by the EndDevice shall not be controlled, even if the import limit is exceeded.
- opModLoadLimW effects a limit on the total active power that can be consumed on a site by controllable DER.
    - Where an EndDevice has been issued an load limit control, only controllable DER that are part of the EndDevice are expected to change their behaviour to respond to the limit. Where a controlled load is consuming active power, this can be offset by local generation

that is part of the EndDevice to allow the load limit to be met. Figure 7 illustrates this by showing the point at which opModLoadLimW applies.

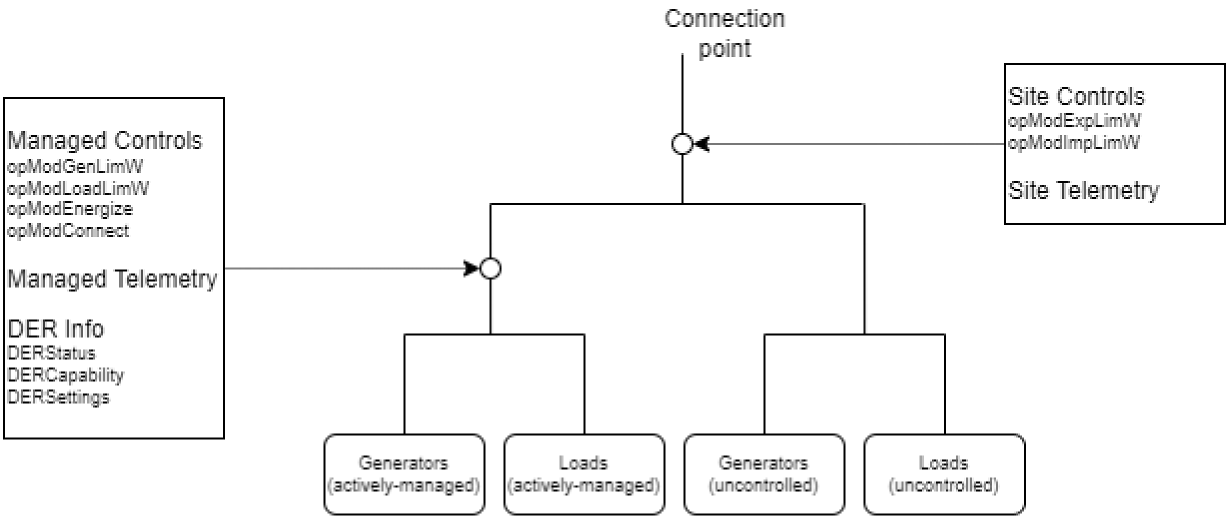

Figure 7: CSIP-AUS site aggregations for controls and telemetry

CSIP-AUS currently defines the ability for a utility server to provide negative limits to a client that the client can then take as an instruction for a minimum level of load (for *opModExpLimW*) or generation (for *opModImpLimW*). While this it supported within the profile it is not an expected feature for clients in general, and it is not tested by the CSIP-AUS test procedures.

6.3 Default Limits

It is important to consider that default export limits apply to a connection point and not an individual DER, and that the fixed import and export limits specified in the connection agreement for that site will apply if the client is unable to register with the Utility Server, or if a DefaultDERControl is not specified by the Utility Server. Where limits specified in a DefaultDERControl are different to the fixed limits applied on site, the default controls specified by the Utility Server should be applied.

For sites with multiple DER, all DER at the site must be installed and configured such that the aggregate response from the DER meets any limits specified. The requirements of the connection standard relevant to the connection point must be considered when determining the appropriate site configuration to meet these fixed limits. Export limits apply to generation-only DER, such as PV. Import limits apply to consumption-only DER, such as EVSE that is capable of charging only. Both import & export limits apply to bi-directional DER, such as BESS and EVSE that is capable of reverse power flow.

FAQ: Why is there no endpoint for DefaultDERControl?

If a DefaultDERControl has not been set for a particular DER then the attribute is not included in the payload. Where a DefaultDERControl is not specified for a DERControlBase, the client shall still implement all relevant active controls. This may occur after server maintenance or patching until DefaultDERControls are re-published for all DERs.

7 CONFLICTING OR OVERLAPPING EVENTS

As per the SEP2 standard, where there are multiple DERControlBase components (e.g., csipaus:opModExpLim, csipaus:opModImpLim) within a DERControl event, unless otherwise specified these shall be considered as independent and are allowed to overlap without superseding. That means an overlapping event will supersede the original event strictly for the DERControlBase specified in that event. The priority of a DERControl is determined by the primacy setting of its containing DERProgram with a lower primacy value indicating higher priority.

Time	00:00 - 00:05	00:05 - 00:10	00:10 - 00:15	00:15 - 00:20	00:20 - 00:25	00:25 - 00:30	00:30 - 00:35
DERProgram B (Primacy #2)		ExpLimW = 10.0 kW ImpLimW = 15.0 kW	ExpLimW = 10.0 kW ImpLimW = 15.0 kW		ExpLimW = 10.0 kW ImpLimW = 15.0 kW		ExpLimW = 10.0 kW ImpLimW = 15.0 kW
DERProgram A (Primacy #1)			ExpLimW = 0.0 kW				GenLimW = 0.0 kW Connect = False
DefaultControl		ExpLimW = 1.5 kW ImpLimW = 4.0 kW					
Fixed Limits	ExpLimW = 1.5 kW ImpLimW = 1.5 kW						
Implied Mode Defaults	GenLimW = ∞ kW LoadLimW = ∞ kW Connect = True Energize = True						
Site Behaviour	ExpLimW = 1.5 kW ImpLimW = 1.5 kW	ExpLimW = 10.0 kW ImpLimW = 15.0 kW	ExpLimW = 0.0 kW ImpLimW = 15.0 kW	ExpLimW = 0.0 kW ImpLimW = 4.0 kW	ExpLimW = 10.0 kW ImpLimW = 15.0 kW	ExpLimW = 1.5 kW ImpLimW = 4.0 kW	ExpLimW = 10.0 kW ImpLimW = 15.0 kW GenLimW = 0.0 kW Connect = False
Notes	Adhere to fixed limits when no controls or default		Export overridden by event with primacy #1 Import of the overlapped event still applies Defaults apply for modes not set			Revert to default limit when no active controls	No consumption or generation allowed due to Connect=False despite still having Exp/Imp Limits

Figure 8: Example of overlapping events and expected behaviour

CSIP (and therefore CSIP-AUS) states that a DER should not resume execution of a lower priority event after completing an overlapping higher priority event. IEEE 2030.5:2023 includes additional details on the expected behaviour under overlapping controls and specifies that the overlapped event shall be resumed on the completion of the superseded control.

The behaviour specified in IEEE 2030.5: 2023 is the preferred behaviour and CSIP-AUS will be updated to move to this approach over time. During this transition period client implementers may implement either approach.

Utility Servers should make best efforts to avoid this scenario by not issuing events that overlap in this manner.

Time	00:00 - 00:05	00:05 - 00:10	00:10 - 00:15	00:15 - 00:20
DERProgram B (Primacy #2)	ExpLimW = 10.0 kW ImpLimW = 15.0 kW			
DERProgram A (Primacy #1)		ExpLimW = 0.0 kW		
DefaultControl	ExpLimW = 1.5 kW ImpLimW = 1.5 kW			
Site Behaviour (Preferred)	ExpLimW = 10.0 kW ImpLimW = 15.0 kW	ExpLimW = 0.0 kW ImpLimW = 15.0 kW	ExpLimW = 0.0 kW ImpLimW = 15.0 kW	ExpLimW = 10.0 kW ImpLimW = 15.0 kW
Site Behaviour (Alternative)				ExpLimW = 1.5 kW ImpLimW = 1.5 kW

Figure 9: Example of resuming previously overlapped events and expected behaviour.

As per SEP2, when a client detects an overlapping event in the same program, the event with the latest creation time will take precedence over the older event. Utility Servers should make best efforts to avoid this scenario by cancelling any superseded events.

<table>
<tr><th>Time</th><th>00:00 - 00:05</th><th>00:05 - 00:10</th><th>00:10 - 00:15</th><th>00:15 - 00:20</th></tr>
<tr><td>DERProgram B
(Primacy #2)</td><td>ExpLimW = 10.0 kW
ImpLimW = 15.0 kW</td><td>ExpLimW = 10.0 kW
ImpLimW = 15.0 kW</td><td>ExpLimW = 10.0 kW
ImpLimW = 15.0 kW</td><td>ExpLimW = 5.0 kW
ImpLimW = 5.0 kW</td></tr>
<tr><td>DERProgram B
(Primacy #2)
more recently created</td><td></td><td>ExpLimW = 5.0 kW
ImpLimW = 5.0 kW</td><td>ExpLimW = 5.0 kW
ImpLimW = 5.0 kW</td><td>ExpLimW = 10.0 kW
ImpLimW = 15.0 kW</td></tr>
<tr><td>DefaultControl</td><td colspan="4">ExpLimW = 1.5 kW
ImpLimW = 1.5 kW</td></tr>
<tr><td>Site Behaviour</td><td>ExpLimW = 10.0 kW
ImpLimW = 15.0 kW</td><td>ExpLimW = 5.0 kW
ImpLimW = 5.0 kW</td><td>ExpLimW = 5.0 kW
ImpLimW = 5.0 kW</td><td>ExpLimW = 10.0 kW
ImpLimW = 15.0 kW</td></tr>
</table>

Figure 10: Example of duplicate events and expected behaviour.

8 SCALING CONSIDERATIONS

As many thousands of SEP2 Clients could be expected to communicate with the Utility Server at the same poll or post rate, consideration should be made when designing clients to handle intermittent communication issues with the server. This scenario is often referred to as the ‘thundering herd problem’.

8.1 Timing of requests

Regular tasks such as sending telemetry at the post rate (e.g., every 5 minutes) could result in the server being overloaded if clients all send the requests at the exact same time (at the top of the hour). Whilst the task needs to be performed at the poll or post rate, a random offset should be applied such that clients do not send all their requests to the server at the same time.

Posting rates should be met within a range of up to 50% of the relevant post or poll rate i.e., if a post rate is set to 60 seconds posting shall occur within 30 seconds of the intended time.

8.2 Retrying Requests

When a server does receive too many simultaneous requests, it will respond with an error such as ‘429 Client Error: Too Many Requests’. Clients should attempt to retry any failed requests after an appropriate delay. Retries should also be attempted for any 5xx Server errors. When determining the delay time, exponential backoff with jitter should be used. The client should make 2 attempts to retry within the polling period, with a minimum of 10 seconds between attempts.

8.3 Telemetry Buffering

If clients are unable to successfully send telemetry back to the server due to communication issues, it should buffer these telemetry readings locally and attempt to resend when communications are re-established. Like storing scheduled events, clients should be able to store at least 6 or more telemetry readings locally, discarding the oldest readings first should the local buffer be exceeded. When sending multiple readings from a buffer, the oldest readings should be sent first.

9 PUBLIC KEY INFRASTRUCTURE (PKI)

Each EndDevice must be issued a unique DER[16]-Encoded X.509 certificate that can be used to securely identify the site or aggregator. The device certificate is used to generate unique identifiers in the form of a LFDI and SFDI via hashing. In the case of cloud proxy connections, a ‘virtual’ LFDI is used.

The device certificate can only be issued via an approved provider, as the certificate must be signed by the Smart Energy Root CA (SERCA).

The certificate file should include the full certificate chain of trust. That is, the file should also include the MICA and SERCA certificates and any MCA that is part of the chain.

9.1 Aggregator certificate structure

Aggregators are able to utilise separate certificates for the base client and the notification server, as described in CSIP-AUS section 5.2.1.2. The notification server shall use either the traditional client certificate or another certificate issued by the utility (or authorised PKI provider) to ensure it is part of the same trust chain.

9.2 Generating a Private Key and CSR

When generating the Private Key, the SEP2 Standard stipulates that the ECC cipher suite SHALL use elliptic curve secp256r1. (This curve is also known as prime256v1 or NIST P-256.)

An example of how to generate this with openSSL is provided below.

```
openssl ecparam -name secp256r1 -genkey -noout -out private_key.pem
```

The private key is sensitive information, ensure that is stored securely and not shared.

Once your private key is generated, you can then generate a Certificate Signing Request (CSR). Note that traditional internet certificates include mandatory fields like the domain name of the website, but for SEP2

[16] Note that in this instance, the abbreviation DER refers to Distinguished Encoding Rules and not Distributed Energy Resource

these must be left blank. To avoid openSSL prompting for these, you need to pass these as empty upfront.

```
openssl req -new -key private_key.pem -out cert_req.csr -sha256 -subj "/CN= /O= "
```

Other methods of generating these files are also available.

9.3 Determining the LFDI

The Long Form Device Identifier (LFDI) is used to identify a site and is calculated based on the device certificate.

The LFDI for a particular certificate can be displayed using the following OpenSSL command:

```
openssl x509 -outform der -in client-cert.pem | sha256sum | head -c 40 | tr '[a-f]' '[A-F]'
```

Or on Windows, the following PowerShell commands:

```
openssl x509 -outform der -in client-cert.pem -out client-cert.der
(Get-FileHash client-cert.der).Hash.Substring(0, 40)
```

9.4 PKI Related Errors

Where an invalid PKI has been presented as part of a request, an “Remote end closed connection without response” error will occur.

9.5 Multiple PKI

Multiple PKI are in use across Australia. Developers should include the ability to transition between different PKI solutions. A target state for Australia is a single national PKI which once established is expected to be the primary mechanism for generating and managing certificates for CSIP-AUS implementations across jurisdictions.

The specifics details in this document relate to the current PKI regimes being used prior to the establishment of the national capability. Once a national PKI is established and is being used by one or more CSIP-AUS implementations, this document will be updated to reflect requirements or capabilities.

10 OEM / TECHNOLOGY PROVIDER APP SUPPORT

The OEM app is what installers use for registration of inverters via in-band registration with the appropriate utility server. In the case of direct to device it can also be used for applying the right certificate.

This section intends to provide guidance for OEMs, Utility Servers and utilities to support features for installer troubleshooting.

10.1 National Metering Identifier

A National Metering Identifier (NMI) is a 10 character identifier for connection points with an additional one digit checksum.

As per the AEMO MSATS NMI Procedure, the checksum is focussed on applications where data entry occurs. OEM apps should require an 11 character NMI to be entered, and validate the checksum. Clients should use the 11 character NMI for in-band registration.

When connecting to a network operator's utility server, it is recommended that clients use the NMI to identify who the network operator is, and map this to the location of the server. The first three characters in a NMI identify the network as per AEMO NMI Allocation List. This can be used to identify which network and hence utility server with which to register.

10.2 Troubleshooting

10.2.1 Failed in-band registration.

For guidance on failed in-band registration scenarios refer to section 4.1.4.

10.2.2 Installation troubleshooting.

A common installation fault is a reversed polarity CT. If the OEM app can identify this prior to testing with utility servers, this will reduce rework. Many utility server provider interfaces do provide information to identify reverse polarity CT, however this is often used by back office staff after installation is complete.