Official Report
Aithisg Oifigeil

DRAFT

Climate Action Committee

Comataidh Gnìomh air a’ Ghnàth-shìde

Thursday 12 November 2026

Session 7

The Scottish Parliament
Pàrlamaid na h-Alba

© Parliamentary copyright. Scottish Parliamentary Corporate Body

Information on the Scottish Parliament's copyright policy can be found on the website—
www.parliament.scot or by contacting Public Information on 0131 348 5000

Thursday 12 November 2026

CONTENTS

	Col.
DECISION ON TAKING BUSINESS IN PRIVATE	1
ENVIRONMENTAL IMPACTS OF HYPERSCALE DATA CENTRES	2

CLIMATE ACTION COMMITTEE
9th Meeting 2026, Session 7

CONVENER

*Stuart McMillan (Inverclyde) (SNP)

DEPUTY CONVENER

*Ariane Burgess (Highlands and Islands) (Green)

COMMITTEE MEMBERS

James Adams (North East Scotland) (Con)
*David Barratt (Cowdenbeath) (SNP)
*Martyn Day (Falkirk East and Linlithgow) (SNP)
*Sanne Dijkstra-Downie (Edinburgh Northern) (LD)
*Donald MacKinnon (Na h-Eileanan an Iar) (Lab)

*attended

THE FOLLOWING ALSO PARTICIPATED:

Dr Gemma Bone Dodds (Lothian Heat)
Simon Kerr (SAV Systems)
Colin Lindsay (Scottish Water Horizons)
Andrea Marinoni (University of Cambridge)
Professor Mark Parsons (EPCC, University of Edinburgh)
Giulio Santori (University of Edinburgh)
Ana Valdivia (University of Oxford)

LOCATION

The Robert Burns Room (CR1)

Scottish Parliament

Climate Action Committee

Thursday 12 November 2026

[The Convener opened the meeting at 09:16]

Decision on Taking Business in Private

The Convener (Stuart McMillan): Good morning, and welcome to the ninth meeting in session 7 of the Climate Action Committee. The first item on our agenda is a decision on taking items 3 and 4 in private. Those items are the consideration of our approach to scrutiny of the Scottish Government's budget for 2027-28, and consideration of our approach to scrutiny of the Scottish Government's improvement plan for Environmental Standards Scotland's air quality investigation into engine idling laws.

Do members agree take items 3 and 4 in private?

Members *indicated agreement*.

Environmental Impacts of Hyperscale Data Centres

09:17

The Convener: The next item on our agenda is our third evidence session as part of the committee's inquiry into the environmental impacts of hyperscale data centres. As I did in previous weeks, I will say a few words about our inquiry before I introduce our witnesses.

The committee agreed its approach to the inquiry on 17 September and 1 October. Our aim is to find out about the environmental impacts of the operation and construction of hyperscale data centres in Scotland. We will consider aspects such as water consumption, air quality, noise pollution and energy use, including back-up generation. We are also interested in the impact of growth in the sector on Scotland's ability to meet its carbon budgets and on the climate change plan. We specifically want to know what systems should be in place to mitigate any negative impacts on the environment and what Scotland can learn from other countries.

We are gathering information in three ways. First, the committee is collaborating with the Economy, Tourism and Energy Committee in running a call for views. I encourage people to share their views on the issue with us. Details can be found on the committee's website.

Secondly, committee members recently visited a data centre in North Lanarkshire. The visit to the DV1 site, which is run by DataVita, allowed members to see how a data centre operates.

Thirdly, we are taking evidence from experts over three meetings. Two weeks ago, we heard evidence from witnesses with an environmental background and then from witnesses with experience of data centre development in Ireland. Last week, we heard from a panel of witnesses from the technology sector and then from a second panel of experts in sustainability and greenhouse gas emissions. Today, we will finish taking evidence on data centres with panels on issues related to water and waste heat.

I re-emphasise that the committee is not engaging in scrutiny of individual planning applications, which are a matter for local authorities. At this stage, we are gathering information that will be shared with the Scottish Government in advance of its production of national planning guidance. We have agreed to consider a second phase of work on data centres and will do that after the planning guidance has been published.

I welcome our first panel of witnesses: Colin Lindsay, developer and hydrogen services manager at Scottish Water Horizons; Giulio Santori, reader at the University of Edinburgh; Professor Mark Parsons, director of EPCC at the University of Edinburgh, who joins us online; and Ana Valdivia, lecturer in artificial intelligence, Government and policy at the University of Oxford. Thank you very much for joining us. We will move straight to questions.

The committee has heard that water consumption by hyperscale data centres is expected to be minimal in terms of the overall demand on a national scale. How can we ensure that the cumulative demand from data centres does not add to regional pressures in areas of Scotland that are at greater risk of water scarcity now or in the future?

Colin Lindsay (Scottish Water Horizons): I think that there are 23 applications at the moment, and Scottish Water is engaged in 15 or 16 discussions with developers. Most of the discussions have been early engagement on what we expect. We have made it clear to the data centre community that, if data centres need tap water, which is obviously a last resort, the systems must be closed loop. A closed-loop system minimises the water consumption impact because it is a top-up system. It might need to be topped up every now and then for maintenance, but that is about it.

On the cumulative impact, part of the challenge is that we have 23 applications and there is a question about how many of these things will actually be built, which is linked to the question of the demand that is out there. However, we will manage that to ensure that we protect existing customer supplies and our water resource. We are trying to influence developers and push them to use alternative water sources such as the treated final effluent from our waste water treatment works. There is an abundant supply of that, which could potentially supply any volume of water, depending on location.

That is the work that Scottish Water is doing to influence where data centres are built and how they are built.

The Convener: Before I bring in other witnesses, I have a question on your written submission, Colin. I was struck by how short it is. What additional research have you undertaken on other countries and territories where data centres are in operation, to help to shape the view of Scottish Water Horizons?

Colin Lindsay: It is probably minimal, if I am honest. That is because the data centre sector has come alive for us only in the past 12 months. We try to engage with the market and the sector, understand it and then influence that view. Twelve months ago, some of the proposals for data centres were suggesting there would be a need for much higher volumes of water than they are suggesting now. That is linked to the work that we have been doing to influence and engage with developers, to ensure that they minimise water demand.

There are horror stories about water disappearing out of town systems because of data centres, but I reassure people that Scottish Water would not allow that to happen. We have controls in place to assess any demand and ensure that there is no impact on customers. If, for whatever reason, there would be an impact—say, for example, that the network would not cope—the developer would have to pay for an upgrade, to ensure that the service was maintained. The developer would pay for that.

The Convener: Ariane Burgess has a quick supplementary question.

Ariane Burgess (Highlands and Islands) (Green): Good morning. Colin, you mentioned encouraging developers that do not have a closed-loop system to move away from using tap water and towards using effluent waste. I recently had a good visit to the waste water treatment plant near Inverness and saw the system there. Would figuring out that technology and how it would work be a requirement on Scottish Water, or would it sit with the developers?

Colin Lindsay: It would be a bit of both. In Scottish Water Horizons, we work with developers, whether they are data centre developers or house builders, to come up with solutions that fit Scottish Water and the developers. We could design the system or they could do it. It is about whatever the right commercial solution is.

The Convener: On my original question, I will bring in Giulio Santori.

Giulio Santori (University of Edinburgh): Thank you, convener and members of the committee. It is an honour to appear before you today to discuss this matter.

On the question of the cumulative demand for water from data centres, water is quite a difficult feature to capture for a data centre, because it comes at the end of the technology and is affected by all the choices that you make upstream. There are choices about the energy input and whether it is renewable, on the back-up system and on the fraction of AI architecture that you are hosting compared to the fraction of high-performance computing. On top of that, we have to distinguish where the water demand will come from—whether it will come from the generation of power or be consumed on site by the cooling system, in

particular. The assumption is that power will be produced from renewable energy and that the standards will set clear indicators on that. The renewable energy factor is an engineering indicator that tells us the fraction of renewable energy with which a data centre is powered. A renewable energy factor of 1 means that a data centre is totally powered by renewable energy, and, if the renewable energy does not have a water demand, the focus moves to on-site water consumption.

With on-site water consumption, you have to make a second choice, which is whether you want AI architecture, and, if so, to what extent—that is, whether you have purely AI architecture, a hybrid that hosts high-performance computing and AI architecture, or exclusively high-performance computing. The trend now is for the hybrid approach; therefore, the question is what fraction of AI you need, because AI architecture is the factor that puts most pressure on water. Because of its power intensity, there is a requirement for cooling technologies that are more advanced than those that are needed for normal high-performance computing architecture.

The problem is the AI fraction that you can allow in a data centre in a specific location. The location will have a certain water availability and water cycle over the year, with especially critical periods in the summer in some locations in Scotland. You need to balance your cooling technology with the water that is available in the location. There are cooling technologies that have close to zero water consumption annually, but there are also technologies that consume a lot of water. The variability in water consumption can be very large. If you go with advanced technology such as direct-to-chip liquid cooling, you get very close to zero water consumption for AI.

There is also the matter of scale—whether the data centre is 10MW, 50MW or 100MW. Whatever the share of AI, the scale affects the amount of water that it needs.

The Convener: I want to come back on a point that you have just made and that you also make in your written submission. In the section entitled "Water quality", you state:

> "Projected increases in water scarcity and drought duration will mandate that hyperscale facilities diversify their water portfolios."

Scotland has more coastline than France, so people find it hard to understand and appreciate a situation in which Scotland could have potential shortages. My question, therefore, is about mandating that hyperscale facilities diversify their water portfolios. How easy is that to do?

09:30

Giulio Santori: As in any business, you do not want to build something that you do not run throughout the whole year. For example, once a business has spent the money to build a desalination plant and is spending money to maintain it, it will want to keep it running for its whole lifetime. That issue affects the price of operating a data centre, so the attitude is that the situation that you describe should be the very last option.

In any case, there are a number of water mitigation actions that can be taken before coming to the installation of such a solution—for example, sourcing water from outside and ensuring flexibility in the management of the AI workload. In my report, I outline a number of mitigation and avoidance measures.

The other point that must be taken into account is that, if you decide to go for a desalination plant, the business will have to use it for more than 25 years, but the industry changes quite rapidly, meaning that what is good right now might not be good tomorrow, in an altered climate. In Scotland, in the current situation, you do not need such plants if you pick your location wisely. At this point, there is no need for desalination. However, in 20 years' time, the technology will have evolved—especially in terms of AI power density—and there will be more frequent droughts in some specific places. If the data centre is located in a place that experiences droughts, that will be an issue in the summer. However, if the data centre is of an appropriate size and possesses internal features that are sustainable in the specific location in which it has been installed, because those issues were taken into account well in advance, there will be no problem.

The Convener: Thank you. Ana Valdivia, could you answer my original question, please?

Ana Valdivia (University of Oxford): Thank you for inviting me. It is an honour to be here, in the Scottish Parliament.

I am an AI expert. I used to train AI algorithms, so I really know how the technology works. However, over the past four years, I have been analysing the impact of the infrastructure of data centres across various geographies, including Mexico, Chile, Spain and the United Kingdom. The first thing that I would like to say is that we need to clarify the idea of the closed-loop cooling system, because the closed loop is actually an open loop. It is not true that the closed loop does not waste water. It wastes water through so-called evaporative cooling, which is what is used in the cooling towers that you see in the data centres, and evaporation is very intense in the summer because of the higher temperatures.

Secondly, of course, not all data centres use water: some data centres also use air. However, air does not have the capacity to take the heat out of the graphics processing unit chips that are used for training AI.

Those are the two things that I want to highlight, because there is a misunderstanding about how the cooling system works. It is a very complex engineering project, so it is important to understand how the different cooling systems work in data centres. The first question that I would ask data centre providers is which cooling system they are going to install in their infrastructure.

The Convener: Thank you. I put the same question to Mark Parsons.

Professor Mark Parsons (EPCC, University of Edinburgh): Thanks for asking me to speak. I am sorry that I am joining remotely, but I am in Boulder, Colorado—it is 2.30 am here.

The Convener: That is fine. Welcome.

Professor Parsons: Thank you.

It is a very complicated area. I run the advanced computing facility data centre, where we have the national supercomputer, and we are building the environment for the next national supercomputer, which will require up to 30MW of power and cooling. Language is really important in this regard, and, in my world of supercomputing, I view AI as just an application of supercomputing.

We view closed-loop cooling as a cooling loop that uses free cooling—it pumps the hot water up on to the roof in a loop through radiators, blows cool Scottish air through it and cools it down in that way. When you run out of heat transfer capability, you chill the water using chillers. Closed-loop water cooling with chillers sounds great because it uses no water at all, but the problem is that chilling uses lots of energy, so it is not really helping the environment, from a power usage perspective. That is why many people have started using adiabatic cooling, as the other respondents have said. Adiabatic cooling basically evaporates water. It is like when you get out of a swimming pool: you are wet, you move your arms through the hot air and you can feel your arm cooling. That is exactly the evaporative approach that is used to cool the water, and there are different ways of doing that.

As the industry moves to using more evaporative cooling, it is really important that we follow standards, such as the ASHRAE standards, including the TC 9.9 technical standards, which tell us how we can use energy-efficient, water-efficient adiabatic cooling. In our opinion, the best cooling to use in Scotland is a mixture of both free cooling—just blowing cold Scottish air through the radiators—and a bit of adiabatic cooling, which uses very little water even on the hottest days. At the moment, our water usage at the ACF is the equivalent of the water used by seven people a day—so, two households' worth. We use a mixture of direct liquid closed-loop cooling, with free cooling, and hybrid cooling, which is a mixture of adiabatic and free cooling.

In Scotland, it is really important that we are very clear about the technologies that we are going to ask hyperscale data centres to use and that we give them targets that they should meet. There are many different ways to do it. The horror stories that we hear about, such as the huge amounts of water being used in the United States, are largely the result of the use of a type of adiabatic cooling called spray cooling, which is extremely water-usage intensive. We do not use that in Europe at all—I do not know of anybody using that.

My message to the committee is that this is a highly technical area and it needs to be approached in that way.

The Convener: In your opinion, therefore, is it fair to say that any examples that we might see in the US are not really comparable with anything that is taking place in Scotland or, more widely, in Europe?

Professor Parsons: The worry is that, without any regulations, a data centre developer could use those very wasteful technologies, which have been used in the US in the past. To be honest, they tend to be used in desert environments, which is why they are so wasteful of water.

Ana Valdivia: I want to mention a risk of using evaporative cooling, which relates to where water that is used in the chillers contains nitrates—for instance, where the ground water has been contaminated with fertilisers. The nitrates do not evaporate, so they accumulate in the cooling system and, when the data centre discharges that water, it creates a lot of problems because the water is highly contaminated with nitrates.

Giulio Santori: I would like to add a variable to the problem. The cooling system that Mark Parsons is talking about might run for only a few days of the year, but those spikes happen when, externally, there is a high ambient temperature and, often, water scarcity—let us say that it is in the summer. If there is a sudden increase in workload, especially for AI, the sudden and near-instantaneous increase in workload can be controlled only by using water for cooling. It might last only for the duration of a heatwave, for example, but, at the peak of that moment, there will be a criticality. Therefore, although the indicator may be excellent throughout the year, in those two days when there is a heatwave there will a peak and the need to use water.

We have just had a lesson on that from the heatwave in late June. What have we learned from that? I looked into four different data centres across the UK. One, in Bristol, is purely an AI data centre; similarly, the one at the University of Cambridge is an AI data centre; there is the supercomputer in Edinburgh, which has a capacity of 5MW or 6MW; and there is the 80MW Google data centre outside London. Of those four data centres, only the one at the University of Cambridge had to be switched off—for nine days—because of a problem with cooling. The one in Bristol was fine because it used direct-to-chip liquid cooling. The one in Edinburgh is a hybrid data centre and it was fine because there were lower temperatures in Edinburgh and because it has a large redundancy in its cooling system. The Google one approached the problem by coping with the inefficiency: an increased amount of the energy consumed was renewable energy.

In my opinion, that tells us a lot. It is proof that, if we do things properly and design a proper cooling system, there should not be a problem. However, we need to plan for the future.

The Convener: Colin Lindsay, do you want to come in on that? I see that you do not.

Mark Parsons, do you want to come in briefly, before we turn to David Barratt's questions?

Professor Parsons: Ana Valdivia has raised the important point that, if data centre owners are to prepare their own water, it needs to be prepared to certain standards. You cannot just use water that you have collected in a pond—it will break the cooling towers, and you do not know what is in the water that you are going to concentrate. I agree with her about that.

David Barratt (Cowdenbeath) (SNP): All the witnesses have already touched on the subject of my question. In particular, I am interested in Mark Parsons's contribution about the technical solutions that exist or that are being developed to reduce water and energy demand. In the context of the data centres that we have been talking about, I heard that the Edinburgh supercomputer uses 5MW. Is that correct?

Professor Parsons: We have a supply of 38MW to the site, but we are using only about 6MW at the moment.

09:45

David Barratt: The data centres that are being proposed are orders of magnitude bigger than that. What are the lessons from the technical solutions that exist or are being developed for data centres of that scale, relative to what we have been talking about so far?

Professor Parsons: We call those organisations hyperscalers, because they take the technology of supercomputing or normal business computing and scale it out, server after server. However, they have a long history of not doing that terribly efficiently—there have been air-cooled servers for many years.

Because of where AI has gone, the chips have huge power requirements. The next generation of chips require 2.5kW, which is more than an iron, and there are thousands of them in rows. As we move to that level of density and power, we really need to consider how much water they will use on very hot days and how much power they will use. There are good ways to do that in an environmentally sensible way that manages the impact on the environment, but you need sensible ground rules, which have to come from people who understand the data centre design world. There are some really good companies in the UK that do that, with lots of detailed modelling.

David Barratt: Thank you. It is over to Giulio.

Giulio Santori: Would you please repeat the question? I am not sure that I captured it.

David Barratt: I asked what technical solutions exist or are being developed to reduce both water and energy demand, particularly at the scale that we are talking about with the hyperscale data centres.

Giulio Santori: I agree with the answer that Mark Parsons gave. The companies could do that through software for flexibility management or through flexibility of workload allocation, for example. The design of their cooling systems could embed cold water storage, to reduce stress and peaks. The solutions exist to cope with the criticalities. To an extent, the solutions will still be avoidance solutions or, in the future, will turn into mitigation solutions. They are still to be discovered and studied. With higher temperatures due to climate change, we will have less water available. However, that can be planned for, because we have models that tell us today what the situation will be tomorrow. The variable that is probably a bit more stochastic is how the AI industry and architecture will develop in the next five or 10 years.

Ana Valdivia: Different solutions are being discovered in regard to cooling systems and how data centres can become more efficient. However, I want to focus on the fact that, when you make something more efficient, that does not mean that it is good. When something becomes more efficient, net demand rises. That is what economists have coined the Jevons paradox. Making all the data centres more efficient does not mean that they will use less energy, because, if everyone is constantly using their computing

power, the net demand will increase. I wanted to bring that to the attention of the committee. We have to look at technical solutions, but we must also see the broad picture and focus on the cumulative effects. We have to look at data centres not as isolated infrastructure but as a sample of different infrastructures.

Regarding energy, I would like to stress something that has been highlighted in the previous sessions, which is that, yes, we have to stop using fossil fuels, but moving into renewable energy and taking it from other industries is also bad for society. We sometimes have to ask data centres to supply their own renewable energy, so that they do not steal that energy from the grid.

David Barratt: Colin, do you have any comments on that?

Colin Lindsay: This is how we would generally deal with things from a water management point of view. Any large industries would always have to have water in storage, to manage their water demand on site and, potentially, so that they could top up at night, when the demand from Scottish Water is not massive. In our main data centres, we normally do several things, as a matter of course, to ensure that any non-domestic customer has the water supply that they need on site—also in case we have a problem on our network.

David Barratt: Picking up on that, what prospects are there of reducing the use of mains drinking water for cooling and of taking water from other sources, such as waste treatment plants or rain? The fact that that water would need treatment has been touched on.

I recently met the British Geological Survey. The Mining Remediation Authority and the BGS have recently undertaken a trial at the UK geoenergy observatory in Glasgow, whereby they injected heated water into an abandoned mine in order to investigate its potential as a thermal energy store—the groundwater is used as a source for cooling and recirculating. Is there scope to expand such work, particularly given that a lot of the proposed data centres would be in former mining areas?

Colin Lindsay: Yes is the simple answer. As I said, Scottish Water is pushing the message that tap water should be the last resort for cooling, and not only in the case of data centres. We are now in 2026, and using tap water for cooling just seems crazy. We need a shift in the cultural mindset in Scotland, because, although everybody thinks that we have lots of water in the country, that is now changing with water scarcity linked to climate change.

It is about how we make that shift. For me, it could be through things such as using treated final effluent, your mine water or whatever other sources could be reused. They would need to be treated to the right level, but there is technology that can do that quite simply, to be fair, so I see it as a massive opportunity.

David Barratt: Instinctively, we are talking about the fact that hyperscale data centres have absolutely enormous footprints. However, if they use a closed-loop system and very small amounts of water, is there not scope for rainwater harvesting combined with storage to address most if not all of the need?

Colin Lindsay: Yes. We have been heavily involved with the hydrogen sector, with which we have been having the same discussion about how it could manage the use of water on its sites, for a start, by holding it for potential use, and about using things such as treated final effluent as a water source. A lot of the hydrogen sector is looking at not using tap water at all but using a different resource, which, to me, is excellent.

David Barratt: Mark Parsons?

Professor Parsons: This is a complicated area. In a typical supercomputer system, cooling is done by having at least three water loops—a loop that circulates around the room and the pumps; perhaps a secondary loop that goes up on to the roof through a heat exchanger, where heat is removed from the water; and then an internal loop in the computers. Quite often, the quality of the water that you have to present to what we call the coolant distribution unit in the supercomputer, which is where the water is managed, has to have very specific characteristics, or the vendor will not allow you to run the supercomputer. That is true of hyperscaler systems as well—they are no different nowadays, as they use the same technologies.

So, if we are to go down the route of data centres not using any tap water at all, there will have to be much more investment in on-site water treatment. You must ask yourselves whether the best way for that water to be treated is by the people who are treating water at great scale and very efficiently or whether you should let everybody build their own water treatment plant at huge expense and almost certainly not terribly efficiently.

It is a complex area, and, as Ana Valdivia said, there are pros and cons in all directions.

David Barratt: Thank you. That was quite thought provoking.

Ana Valdivia: To add to what Professor Mark Parsons just said, we can sometimes look into having technological innovations in data centres and ask that there be a treatment plant to clean the water that is discharged. The solutions are sometimes already there.

David Barratt: Are there any further comments on that point?

Giulio Santori: Your question made me think about the committee's request to express some thoughts about the situations in other jurisdictions. I did a bit of research into the situation in Lombardy, which recently changed its legislation to forbid the utilisation of water at all, except in exceptional circumstances, and to allow only the use of reclaimed water by data centres. I provide that example to add an element to the discussion. I am not sure whether it is the right way to go, but it is an approach that has been used in places that are even warmer than Scotland in the summer.

Sanne Dijkstra-Downie (Edinburgh Northern) (LD): I will introduce the theme of digital electronic waste from hyperscale data centres, and my first question is about the lifetimes of the different components that are used in data centres. We discussed that issue with witnesses last week, who told us that, whereas the hardware of ordinary data centres can last six, seven or eight years, the lifespan could come down to perhaps two or three years in the case of AI-related data centres. Mark Parsons, what is your assessment of the life cycle or lifespan of that kind of technology? Are we talking about technology breaking down or components becoming obsolete because of advances in technology?

Professor Parsons: When people buy a new system—I am talking about when people buy thousands of GPUs or central processing units—they tend to find that there are lots of failures in the first six months and that the number of failures then gradually decreases. Those failures will have already been present when the device was manufactured. There is then quite a long period—probably four, five or six years—in which there are far fewer failures, and then the number of failures normally starts to rise again at about year 5 or 6.

When we had just CPU-based computers—the old-fashioned supercomputers—they could be run quite happily for six or seven years. However, the advent of a particular type of memory in GPUs called high-bandwidth memory has resulted in a more finite lifespan. For example, I understand that NVIDIA nowadays warrants its GPUs for only five years because the number of failures increases so rapidly after that point. We have seen that happen ourselves with our limited number of GPUs.

I do not agree that people will use systems for only a couple of years. Instead, the latest GPUs are charged with expensive pricing for the first 18 months to two years while the older ones are charged at much lower rates—we see that happening with the hyperscale offerings. Therefore, when people buy time to use GPUs, the clever people are using the ones that are three or four years old, because they are much cheaper. In the market today, H100s and H200s—and, in places such as Microsoft, V100s—that are five, six or seven years old are still being sold as part of the offering from companies. The lifespan of most of those technologies is currently about five or six years.

Sanne Dijkstra-Downie: Thank you. As part of the evidence that we heard last week, the point was made that, although some of those pieces of equipment can be reused, there will still be a massive overcapacity of that type of technology, simply because of the amount of computing that is required across the country. Will you reflect on that point?

Professor Parsons: We are now in a different world. Let me talk about the situation five years ago, before the AI bubble happened.

Supercomputers were always the largest computers in a country, and, at the end of a supercomputer's life, you normally give the disposal rights to the vendor, because they are the experts in disposal. They also charge you less at the start if they hold the disposal rights, because they can then choose how they recycle or resell the systems. For example, one of our systems was sold to an industrial organisation and, happily, it had an extra life for three or four years after it was decommissioned as a national supercomputer.

Nowadays, there will be a huge amount of GPUs coming out, with all the wires, all the memory and all the other bits and bobs, and the big companies of this world—the HPEs and Bulls—and the hyperscalers themselves have recycling organisations. One of the largest in Europe—HPE's European headquarters for recycling—is, in fact, in Erskine, at the old Compaq plant. If the committee has never been there, I recommend that you go and have a look around. The people there will explain how the industry works, and it is a very interesting bit of the industry.

Sanne Dijkstra-Downie: Great. Thank you.

Ana Valdivia, will you reflect on what you see as the lifespan of that kind of equipment and share any reflections on the reuse or recycling of that technology?

10:00

Ana Valdivia: I have done fieldwork on e-waste, and I realise that people do not usually look at the state of the components. Unfortunately, I have a different story to that of Professor Parsons. I have been told by industry providers—that is, by people who work with e-waste in the UK—that the lifespan of a GPU is one year. When we are training large language models, the AI workload is so intense that the chip dies after one year. If you read the

annual reports of the big tech companies, you will see that the official numbers say that the lifespan is between three and six years. However, if you talk to people who actually deal with e-waste, they will tell you that they are seeing GPUs that are only one year old.

What happens at the end of the lifecycle of a GPU is also a very interesting story, because not a lot of people know what happens when a data centre decommissions GPUs. I have been analysing e-waste flows in the UK, and 80 per cent of the e-waste is sent to Belgium. You might wonder why everything is sent to Belgium; it is because there is a recycling plant there, which is where most of our e-waste ends up. Unfortunately, recycling is also a dirty business, and what happens there is creating a lot of environmental impacts for the people who live near that recycling plant.

Sanne Dijkstra-Downie: I will bring Professor Parsons back in, because he gave a different assessment of that and I saw him shaking his head at the mention of a one-year lifespan. Are you talking about different types of technology?

Professor Parsons: No. We drive GPUs as hard as somebody doing big LLM training.

It is absolutely true that GPUs break three months or a week after you buy them, but it is a numbers game. As I said, you see quite high attrition rates at the start, then things level off for three or four years, and then you see the attrition rates growing again. Yes, GPUs die after the first year, but it is not a large number of them. It is definitely a higher number than in the old CPU world, though—everybody in the industry knows that.

Sanne Dijkstra-Downie: Giulio Santori, do you have anything to add?

Giulio Santori: I cannot comment, as I have not done research on that.

Sanne Dijkstra-Downie: That is fine. Colin, I do not think that it is your area either.

Colin Lindsay: Not at all.

Sanne Dijkstra-Downie: That is absolutely fine.

Another issue that we discussed last week is embodied carbon in the production of the hardware—the emissions associated with the manufacturing and also the water use in the manufacturing of this kind of technology. Colin, do you have anything to say about what we should be thinking about here?

Colin Lindsay: As with any aspect on the water side of things, it is simply about reusing and recycling water, if we can. For any process that uses water, we need to ask whether it needs to be tap water. Tap water comes with a carbon price based on the fact that it has to be treated and pumped. However, if you need to do that, it is about considering the most efficient way of treating the water.

Sanne Dijkstra-Downie: Do any other members of the panel have anything to add specifically on the topic of embodied carbon in the production of the hardware?

Ana Valdivia: I can comment on water usage. At the University of Oxford, we have been analysing the water consumption not only of data centres but across the AI supply chain, from mineral extraction to waste. We found that the semiconductor industries are more water intensive than data centres. In analysing water usage, we found that the water consumption of big-tech companies such as the Taiwan Semiconductor Manufacturing Company—TSMC—which manufactures GPUs, is much higher. It is interesting to see how different stages of the AI supply chain have different levels of water consumption. I would say that the semiconductor industry is more water intensive than other industries.

Professor Parsons: I completely agree with Ana Valdivia on that.

This is a problematic area for the industry. People in the big companies that work in this area say that they find it very hard to track the embodied carbon all the way back to the mining of the materials from which the chips are produced, wherever they are found. We ask them to tell us how much embodied carbon there is in a server, but the numbers that come back from procurements are so wildly different that we cannot use them as an indicator, because we just do not think that the numbers make sense. Governments should be pressing anybody who sells servers to get a much better handle on the embodied carbon and on the damage that they do to the environment.

I completely agree on the point about the manufacturers, such as TSMC, that make the chips. They use some really nasty chemicals to do that. Such companies manage those chemicals well, but a huge amount of pollution is still created.

Sanne Dijkstra-Downie: That leads me to my final question. Should we be considering any other forms of waste from the operation of hyperscale data centres?

Professor Parsons: I have been making myself unpopular in the data centre community in Scotland by saying that we should follow the Scottish Government's guidelines on embodied carbon in the development and building of buildings. At the moment, we are working within

the 2030 target in the developments that we are doing. It is important to take that target into account when such developments are happening.

Ana Valdivia: We should also focus on the toxic waste that is generated. We talk about energy and water but not about the toxic waste that data centres generate.

The Convener: Has the discussion on embodied carbon been taking place in an international space, too? Has the matter been raised at a United Nations climate change conference of the parties—COP—present or past?

Professor Parsons: I do not know. The last time the Scottish Government published information on that was around the time of the COP that we held here, in 2021 or 2022. That is not formal legislation; it is work that was done by the Energy Saving Trust, if I remember rightly—or by one of those organisations.

The Convener: I ask that question because the tech industry is a global industry, as we all recognise, and parts will come from different countries—it is not as if everything comes from just one place—so there needs to be a wider global emphasis on that issue.

Martyn Day (Falkirk East and Linlithgow) (SNP): In previous weeks, we have heard evidence about the Scottish circular economy strategy, which was praised as being far ahead of policy in other parts of the European Union. What would you consider to be an effective circular economy policy, and what should that look like for data centres?

Professor Parsons: We need to separate data centres from the systems that they hold. My own data centre was originally built in the 1970s, and my oldest room is 50 years old now. It has grown up over time: it has four rooms now, each slightly different. Within their lifetime, those rooms have held many different systems, so we need to think about the life cycle of the data centre as being separate from the life cycle of the many systems that it will hold over its lifetime.

We have conflated hyperscale data centres with the AI technology that they will potentially contain, but those are actually two different conversations. One of the most wasteful things that we could do is allow the hyperscalers to build a data centre, fill it with equipment, run it for five or 10 years and then knock the thing down. That is not an environmentally sensible way to behave. Yes, you upgrade the rooms over time. For example, we are replacing the cooling loop that has been in place for the previous two national supercomputers because it is almost 20 years old and has been on continuously for 20 years. That is fine, but we are not rebuilding the building—we are reusing all of the infrastructure in it.

Martyn Day: Giulio Santori?

Giulio Santori: This is not my area of research.

Martyn Day: No problem. Ana Valdivia?

Ana Valdivia: The first measure that I would ask for in the circular economy context is an extension to the lifespan of chips. I would ask not just the data centres but the semiconductor industry to focus on making the chips more reliable, so that we do not have to decommission them and create more e-waste every one, three or six years.

I would also like to draw attention to the ingredients that are embedded in a GPU. The mineral that is most commonly found in a GPU is copper: a GPU is 92 per cent copper; the next most common element is iron, and the next is chromium. The copper can be extracted and reinserted into the economy, but the other minerals cannot. Therefore, when you try to put the GPUs into the circular economy, they are not valid because of the minerals that are embedded in them. The minerals have to be very pure, because these are very sophisticated chips, so the minerals that can be extracted and recycled cannot be inserted into a new GPU—brand new minerals have to be used.

Martyn Day: I was looking at some of the components in these processors and wondering about asking a further question, which I will ask now since you have raised that issue. How do we square the circle with regard to things such as lithium in batteries, cobalt and helium? We do not even produce helium in this country.

Ana Valdivia: The solution is the circular economy. Lithium can be recycled, but there are other minerals that cannot really be recycled. I have been told by e-waste recycling providers that the circular economy is useful for minerals such as gold or silver, because they are very valuable in the market, and the price of copper is increasing, year by year, because it is more difficult to extract. Unfortunately, however, not all of the metals in a GPU can be recycled.

Martyn Day: Colin Lindsay, do you have any thoughts on that?

Colin Lindsay: From a water point of view, and linked to the circular economy, one of the things that we are exploring just now with smaller companies is using treated final effluent for open-loop technology in order to reduce the energy consumption of small data centres. My understanding is that about 20 per cent of the energy consumption is for cooling. The companies are building smaller sites and are looking to make them more efficient, and using treated final effluent

as the water source for cooling would potentially allow them to use an open-loop system. However, that work is in the very early stages. Off the back of that, we need to understand what open loop means. What comes back? Does anything come back from the open loop technology? That work is in the very early stages, but that might present an opportunity as well.

Martyn Day: My last point goes back to some of the questions that other members asked earlier, and it relates to the circular economy and e-waste. What sort of regulations should we be looking at to prevent waste from our data centres being exported to other countries that have lower regulations? Ana Valdivia, would you like to start?

Ana Valdivia: The Basel convention, which was implemented in the 1990s, tells countries not to export e-waste to other countries. I analysed United Nations data on e-waste flows and discovered that, for instance, when e-waste contains valuable materials that can be recycled, it stays in the European Union. For example, 80 per cent of UK e-waste ends up in Belgium. However, interestingly, when that e-waste does not contain those metals that are useful in the circular economy, it is shipped to third countries, such as Indonesia.

10:15

Martyn Day: That is very interesting. Thank you.

Donald MacKinnon (Na h-Eileanan an Iar) (Lab): I want to turn to the siting of hyperscale data centres. I would like the panel to consider what a responsible one looks like and what criteria should be used to decide where they are located. It would be helpful if, in answering the question, the witnesses could consider what they understand the term “hyperscale data centre” actually to mean.

Colin Lindsay: Location is really key from our point of view. If we could influence where data centres are built, they would, for a start, probably be in areas that are not water scarce. We would then look at whether they could be sited next to large waste water treatment works, because that would mean that they could use the treated final effluent, rather than tap water, as a water source. We would also consider which other water sources could be used from a sustainability point of view. From a water perspective, location is key, but the way that the system currently works is that the grid absolutely drives where the data centres are being built. Water is a secondary aspect of it.

Donald MacKinnon: What is your understanding of the term “hyperscale” in relation to data centres?

Colin Lindsay: My understanding is that a hyperscale data centre is probably a large data centre that is more than 500MW in size.

Ana Valdivia: I would like to raise the Aarhus convention, which is a UN environmental treaty that the UK signed in 2005 and that came into effect in Scotland in 2008. The convention establishes three core procedural rights regarding the environment: access to environmental information, public participation in environmental decision making, and access to justice to challenge decisions. A responsible data centre should comply with the Aarhus convention, ask local communities what they think about projects and give them information about how much water and energy it will use, so that those communities can also participate in decision making on the plan.

Donald MacKinnon: That was really helpful, as it gave a grounding in the international context as well, and the relationship with the local community and how it interacts with the project. Do you have any thoughts on what the term “hyperscale” means in relation to those rights?

Ana Valdivia: There is actually a categorisation for data centres. You might know that there are eight data centres—or co-location data centres—in Scotland. The new category considers them to be hyperscale when they use more than 100MW. I consider such data centres to be hyperscale.

What is interesting about a hyperscale data centre is that it can tell you, “I’m going to use 100MW this year,” and then come to you two years later and say, “I’m now going to scale up and use 500MW.” That feature is intrinsic to hyperscale data centres: they can scale up without limit in terms of energy capacity.

Giulio Santori: You used the key word “responsible” in your question, which embeds the green dimension and goes beyond it. I read the term “responsible” as a synonym for “sustainable”, and sustainability has three dimensions: the social aspect, on which I am not an expert; the environmental aspect—I can talk about the nexus between cooling and water; and the economic aspect, because data centres must be good in terms of the money that they make. A balanced approach means considering all three dimensions at the same time.

What is a responsible data centre when it comes to on-site water demand? I would say that a responsible data centre is one that takes into account the local context of the place where it is installed, including local availability of water and competition with other users of water in the same location. That is the downstream context. Upstream, there is the question of the availability of renewable energy in the location in which the data centre is installed.

We are already stepping into the definition of a green data centre, which, in my opinion, should take into account the five categories of engineering standards around data centres: power usage effectiveness, the renewable energy factor, the energy reuse factor, the cooling efficiency ratio, and water usage effectiveness. Those are five parameters that describe the performance of a data centre. A responsible data centre is a data centre whose performance is optimal across those five dimensions.

Donald MacKinnon: Thank you for that. What is your understanding of the term "hyperscale"?

Giulio Santori: I would not use the word "hyperscale" at the moment, because the thresholds move. The power density of a rack is projected to scale up to 1MW in the next 10 years, so you might end up with edge data centres, which are small data centres that need a lot of power. I recommend that you get rid of a classification that is based mainly on the number of racks and rely on a classification that is based on the power that is consumed, especially the power that is consumed by the AI fraction in the data centre.

Professor Parsons: I do not think that the question of where we should site big data centres is particularly difficult. We should use land that has already been designated as industrial land, which has been deindustrialised as steel plants and power plants have disappeared over time. Often, those places already have good access to water and electricity. I do not think that we should build big data centres in rural locations, which would harm the Scottish landscape. There are more than enough places in Scotland where we could build very large structures that would have little impact on the environment and on people's lives and that would—if we do it correctly—provide jobs and bring wealth to those regions. I think that it is really a planning matter and that we need some planning guidelines. Those should be produced as quickly as possible.

In my world, we do not use the term "hyperscale". We refer to a group of American companies as hyperscalers, which we do with a bit of derision, to be honest, because they have always pooh-poohed the world of supercomputing, believing that they do it better. In fact, they are now in our world, and they are discovering all the challenges of running such large technologies, with all the water cooling, power distribution, management of fire suppression and so on that that involves.

If you want to use the word "hyperscale", you should think of a hyperscale data centre as a very large data centre. I talk about our data centre rooms. The new one that we are building is a very large 30MW room that is about the size of a medium-sized Lidl. If somebody asked me, "What does a hyperscale data centre look like?", I would say that a hyperscale data centre is on the same sort of scale as the big Amazon warehouses that you go past when you drive along the motorway up to Perth.

As Ana Valdivia said, in some places in the world, they scale out hugely, at 400MW, 500MW or more. One of the beauties of putting them on bits of land that we have already remediated—pre-industrial land—is that there is often a lot of such land. Again, if planning works, it would be saying, "This is what we're building today and what might be built in the future. Does it make sense to set aside this land? Is there power, water and cooling? Are there the people to do it?" It becomes a planning issue.

We have recently been involved in lots of conversations with power companies and providers who have visited our data centre, and, in truth, I think that there will only ever be two or three very large data centre sites in Scotland. Apparently, about 25 proposals are on the waiting list for power at the moment, but I do not think that that will happen—I think that we will have a small number of very large sites in Scotland.

Donald MacKinnon: Thank you, Mark. Those were helpful insights.

I want to pick up on your point that, ultimately, this leads into the planning guidance. The Scottish Government has committed to publish planning and spatial guidance at the end of the year. What would panel members like to see in that guidance? What is important?

Giulio Santori: I would like to see clear maximum and minimum thresholds across the five categories that I mentioned earlier.

Colin Lindsay: Scottish Water has already given input to the guidance. Our view is that if tap water is going to be used, it should be used in a closed-loop system, but that tap water should be avoided as much as possible.

Donald MacKinnon: You want very clear guidance on that.

Colin Lindsay: Yes.

Ana Valdivia: I would add that we should have more transparency. We need to make the publication of environmental reports obligatory for the industry, and communities must be involved in decision making about data centres. If I were an Edinburgh citizen and my house was going to be affected by a data centre being built nearby, I would want to know more about it, have a say and be part of the decision making. That is one of our demands on the industry.

Donald MacKinnon: Mark, did you have something to add?

Professor Parsons: I will not name the person, but a director of planning at one of the big councils in Scotland visited the data centre a couple of weeks ago and was surprised by how quiet the data centre was when it was running. We need to understand correctly the urban and rural environmental impacts, and what data centres really sound and look like. There is some unhelpful publicity about data centres at the moment, which is making people scared of them when it is not necessarily the case that they should be scared. I have managed to get planning permission for our rooms over the years largely because we are tucked away and, for the past 20 years, have been on what has already been designated as a science zone—the Midlothian Science Zone. As a result of that zoning and the fact that we are not near houses or whatever, we have managed fine, to be quite honest. It really comes down to proper planning on where we want to put data centres.

Donald MacKinnon: My final question is on the UK Government's approach to AI data centres and, in the Scottish context, the AI growth zone that has been identified in North Lanarkshire. How helpful is it to identify growth zones and, effectively, concentrate development in one area?

Ana Valdivia: I would answer that question with another, which is how many data centres do we need? How much data are we generating? I raise those questions with policy makers and nobody knows.

Donald MacKinnon: That is a question that the committee is very interested in drilling further into.

10:30

Professor Parsons: We could also think about turning that question around. In thinking about where we want to put data centres in the UK, we should think about where the best place to do that for the UK would be.

There are carbon intensity websites, which I look at quite a lot, so we could look at that aspect. In general, Scotland is at 1g or 2g of CO_2, whereas the south-east of England is at 250g or whatever. There is, therefore, very little sense in building more and more data centres down in the south of England. Scotland has power and it has space, and it needs to create a vibrant growing economy. We should not allow ourselves to think that data centres are bad—if we want to have AI and places to store our music and our photographs, as well as email and Microsoft documents, there need to be these things called data centres. Scotland should play its part, but it should do so responsibly and environmentally.

Giulio Santori: I can see a benefit from the AI growth zone, as it is pushing the sector. The reason that we are talking intensively about data centres is thanks to the AI growth zone. The future is going to be even worse than now, so it is better that we get on with addressing the topic right now, because, after the GPUs, we are probably going to see quantum computing, for example. There are already some commercial quantum computers—the numbers are small, but they already exist. We need to be prepared.

With regard to how much we need, the question, in my opinion, is, what is the demand from the UK population? The demand for AI and digital services is huge, and, if you do not meet the demand within the boundary of your country, you are going to do so outside the boundary.

Donald MacKinnon: There are no further questions from me, convener.

The Convener: Before I bring in Ariane Burgess, I will go back to Giulio Santori's final point. Does that indicate that data sovereignty is important with regard to data centres?

Giulio Santori: Yes—it was a clear reference to that.

The Convener: Thank you. I see that Mark Parsons is nodding in agreement.

As Ana Valdivia and Colin Lindsay have no comments on that area, we will move to questions from Ariane Burgess.

Ariane Burgess: I have a number of questions. I will, in the first instance, direct each question to Ana, and then I will invite the other witnesses to indicate if they want to come in.

I am interested in getting a sense of whether, in the context of Scotland's climate targets, you think that Scotland would be able to host hyperscale data centres. We have heard that there are some data centre applications in the pipeline, although we might not necessarily get the 23, 24 or 25—the number keeps changing—that we have heard about. Nonetheless, if we were to accept all the applications that are currently in the planning system, or even two or three of them, would we remain on target to meet our 2045 climate targets?

Ana Valdivia: I cannot really comment on Scotland's climate targets, but I can bring in another example. We are analysing the carbon emissions relating to the usage of LLMs at UK universities. We are asking all the Russell group universities—there are 24 institutions—for information such as how many tokens they are generating. In the worst-case scenario, which is that everyone is using AI agents, which demand a lot of tokens, with a very intensive workload, the carbon emissions from LLMs are 80 per cent of the

emissions from air travel. As you know, universities have to report when academics travel and the carbon emissions related to that.

I cannot comment on the specific question of Scotland's climate targets, but there we have an example that shows that AI is becoming an important and significant source of carbon emissions in the university context.

Ariane Burgess: Does anyone else want to come in on that specific question?

Giulio Santori: My expertise is not in the embodied carbon aspect of data centres, but the question resonated with me in relation to the issue of back-up generators, which are activated in exceptional circumstances that, nowadays, are related to problems with the grid. If a data centre produces its own energy from renewable sources, the emissions and the water usage that are related to the back-up generator must be reinterpreted, because, if they involve renewable sources on site, we have solved another problem.

Ariane Burgess: Thank you for highlighting that—it is very helpful.

We have been talking about the planning system and the guidelines that the Government wants to introduce. However, looking to the future, I am interested in public reporting. The EU now requires large data centres to report annually on energy efficiency, renewable energy use, heat recovery and water consumption. Can you give us your reflections on that? Ana Valdivia, in Scotland, should developers and operators be required to publish information on things such as embodied carbon, e-waste, back-up generation, noise, biodiversity impacts and waste heat utilisation? If so, which metrics would provide the most meaningful public accountability? You and Giulio Santori have both reeled off a few of those elements.

Ana Valdivia: Yes, developers and operators should report those metrics. Scotland should also look at article 12 of the European Commission's energy efficiency directive, which asks data centres above 50MW to report those environmental metrics, which means that data centres below that threshold do not have to do so. I recommend that Scotland create a publicly available database that provides information about every provider in Scotland in terms of how much energy they demand, how much water they use and so on.

That regulation focuses on data centres that are above 50MW, but you have to think about what happens when you have, for example, three providers that build three data centres close together that come to 40MW each, which is something that I saw in Mexico. If you put in place such a threshold, people will develop data centres that are below it but they will build all of them together, so we have to look at the cumulative effect as well. If you put in place such a regulation, you should bear that issue in mind. In Spain, the royal decree on data centres has addressed the issue by saying that the energy demand of all the data centres within a certain distance of each other must be taken into account.

Ariane Burgess: Should those metrics be independently verified in the same way that environmental reporting is audited in other sectors, rather than our relying on the operator to self-report?

Ana Valdivia: Yes, there must be independent accountability. Giulio Santori might want to say more about this, but I would stress that some of those metrics have limitations. For example, the power usage effectiveness metric does not really tell you how efficient a data centre is; it gives you only a broad picture. You would have to ask the operator whether, for example, it shuts down the servers at night if they are not being used. That is an efficient thing to do, but it is not reflected in the PUE metric.

Biodiversity is an important factor in terms of the environmental impacts of hyperscale data centres. It is important that we do not just look at nature in terms of metrics, but also make observations. For example, if a data centre comes to an area where birds are nesting, it will be important to ask experts in biodiversity to observe how the data centre operations impact on those species.

Giulio Santori: I recently watched a similar session that involved colleagues in the USA, and they said that there is a problem with the transparency of data. We need to monitor performance across the different indicators, as the Scottish Environment Protection Agency already does in many other industrial sectors. Monitoring is very important, as is transparency across all the data. At the moment, there is not much of that.

Under the standards, the metrics that I listed are normally averaged over the year. We do not want that approach, because the problems might exist in the peaks. Rather, we want to have thresholds that are maximums or minimums, depending on the indicator. The granularity of the monitored data is also important.

Professor Parsons: I agree with all of those points. Also, we must not miss the fact that there is often, although not always, a split between the data centre owner and the tenant. There should be a requirement on the data centre owner to work with their tenant to minimise electricity and water use. If we do not have that requirement, there will be no impetus for either side to spend any time worrying about that. We can get them to publish

their numbers, but, if the data centre owner cannot influence the tenant, there is nothing that the owner can do about the numbers.

Ariane Burgess: That is a really helpful point, because I have been puzzling over that. The developer will not necessarily be the person who will actually operate the site. It seems to me that they can provide just broad-brush designs for what will be there in order to get it through planning, and we will not get the detail until something is built out and actually operating.

Ana, I ask you to respond first to my next question, because you have named a lot of other countries as you have been talking. We are interested in examples of hyperscale data centre policy development in other countries, and I am interested in hearing what you believe will be relevant for our forthcoming Scottish guidance. You talked about the Spanish royal decree, and you also mentioned things that are happening in Mexico, where smaller sites are being combined and are getting under the wire for environmental impact reporting. Do you have any helpful suggestions of countries that we should look at? We are already aware of some examples, such as the guidelines that came out in the summer in Lombardy.

Ana Valdivia: The Spanish royal decree has not been accepted yet, so it has not been enforced. I think that it has been mentioned in previous evidence to the committee. Basically, it imposes on data centres a duty to buy power purchase agreements. They demand that data centres use a certain amount of renewable energy, although there has been a lot of back and forth on that, because communities also see negative aspects in renewable energy. Spain is a good case study in that regard, because renewable energy has had social and environmental impacts there, which are sometimes not mentioned or underscored.

Another interesting aspect of the royal decree is that it requires the monitoring of data centres. If a provider says that their centre will use a certain amount of energy but it actually uses more, the royal decree will impose fines on the provider. It can even reach a point where they will be disconnected from the grid if they do not do what they said they would. That is an interesting measure. In other jurisdictions such as the US, data centres have said that they would use a certain amount of water but it has turned out that they have used more.

My other point relates to environmental impact. Under the Aarhus convention, which Scotland is signed up to, we have to guarantee that communities have access to information and are part of the decision making. In some jurisdictions, data centres are not really providing information to the public.

Ariane Burgess: Okay. Thanks. As no one else want to come in with suggestions of other countries that we could look at for our guidelines, I will move on to a question about e-waste.

As was mentioned earlier, NVIDIA has only a five-year warranty, and the time period can be even less. We are tied into the problem of built-in obsolescence. Most of us will know the history of the light bulb, for which that was a conscious act. Is there an international convention or something whereby we could require technology companies to ensure that their products have more longevity?

Ana Valdivia: The European Commission has a specific framework on that. I cannot remember its name, but there is a legal policy on the obsolescence of appliances. It applies to washing machines and laptops, but unfortunately it does not apply to chips.

You could ask whether NVIDIA could expand the lifespan of its GPUs. Something that has not been mentioned is that NVIDIA has a monopoly. When AI systems are trained using GPUs, there is software called CUDA—compute unified device architecture—and NVIDIA has a monopoly on that. Without that software, we cannot really train AI systems. If NVIDIA decided tomorrow that CUDA could no longer be used with A100 GPUs, they would have to be decommissioned, which would create a lot of waste. NVIDIA therefore has a lot of power over e-waste flows, because it can decide when its software will no longer work with its hardware.

Mark Parsons might want to comment on that.

Professor Parsons: NVIDIA has been extremely successful with its marketing. In truth, the GPUs of its competitor, AMD, are just as powerful, but NVIDIA has a very successful software stack. We are in a different world nowadays, and we are buying both AMD and NVIDIA GPUs, as well as other devices. The market is becoming more competitive.

On the previous point, it is not really obsolescence. In the case of not just NVIDIA but all the manufacturers, you can think of chips being like roads that get narrower and narrower. They are now down to 2 nanometres, which is tiny. In effect, the chips are wearing out as they get hot. They suffer thermal damage and, over time, the memory fails. There is nothing that we can do about that; it is just physics. There is a drive for hotter and hotter chips that can do more processing more quickly to provide what we need, particularly for AI and for other bits of supercomputing. That simple fact is driving us to

the edge of physics, and there is nothing that we can do about it whereby the chips will survive.

Ariane Burgess: Since our meeting last week, we have had a big change in the landscape in the United States, with the midterm elections. It may not be appropriate to ask you this question, but I wonder whether you have observed a shift in tone on AI in general in the US, which has been driving its development globally.

Ana Valdivia: I have been following the news regarding the elections in the US and how the data centre industry impacted on votes. Some results point to the fact that neighbourhoods near data centres have not voted for the current Government. We are seeing data centres becoming political assets, because people are not convinced about the benefits that the industry will bring to their neighbourhoods.

The Convener: With that, I thank our witnesses for their evidence. I will suspend the meeting for a few moments to allow a change of panels.

10:50

Meeting suspended.

10:59

On resuming—

The Convener: I welcome to the meeting our second panel of witnesses. Before I introduce them, I put on record that James Adams has sent his apologies.

Before us are Dr Gemma Bone Dodds, chair of Lothian Heat; and Simon Kerr, head of heat networks at SAV Systems. We have received apologies from Andrea Marinoni, who is ill and could not be with us today

We will move straight to questions. How feasible is it for the waste heat from data centres to be captured and used in heat networks? I put that question to Gemma Bone Dodds first.

Dr Gemma Bone Dodds (Lothian Heat): It is very feasible. There are data centres, particularly in Denmark that already utilise the heat from the processors. Our colleague Viegand Maagøe has been involved with lots of planning applications in Denmark, as well as with numerous proposals for connecting upcoming large-scale data centres to existing heat transmission pipes, to deliver heat to the city centres, where it is needed.

On the technical side of things, Simon Kerr will probably give you a better answer, but it is worth noting that closed-loop water cooling systems are much better, not only for water use but for heat capture. Such systems are the easiest to install, but the best approach is to plan for a data centre to reuse its heat from the start. Retrofitting an existing facility will always be much harder to do.

Simon Kerr (SAV Systems): Heat transfer is a simple technology. Our fridges, air conditioning, heating and even gas boilers all operate on the basis of transferring heat. Managing heat transfer is very simple when you insist on it at the planning stage.

To get planning approval, energy-from-waste plants in Scotland must have heat transfer and heat off-take infrastructure built in. Data centres that are built with heat transfer in mind should be required to do that, too. Planning authorities should insist that the heat exchanger infrastructure is built in to capture the heat and reuse it in heat networks.

At the minute, heat networks are struggling to get off the ground for various reasons, and we have been flatlining for 20-odd years. I do not know whether anyone has seen the graph for heat network growth from 2030 onwards—it is practically a horizontal line up to 2045. Achieving such growth is not realistic.

We need to start getting serious about it. We are at a crossroads. The data centres have been handed to us like a gift from heaven, if we get it right. We need to get the data centres, the energy-from-waste plants and the brewing industry, as well as the national health service laundry services and all the other stuff that is continually losing heat, to understand the value of that heat. We also need to put in place legislation, including in planning, so that anyone who intends to discard heat has to build in heat transfer systems that enable heat networks to access it. Otherwise, we will still be flatlining in another 20 years.

As I said, the solution is simple to implement, but data centres and others will not do it on their own. We need legislation.

The Convener: Dr Bone Dodds, you mentioned retrofitting, which can be a lot more expensive. In Denmark and in other Scandinavian countries, heat networks are very common, in contrast to here, and they are built in at an early stage. There are 23 to 24 data centre applications across Scotland, but it is widely considered that the vast majority of them are speculative and will not go forward. The evidence that we heard last week suggested that we could have more but smaller data centres rather than fewer but larger data centres. What would be the best outcome for heat networks if any of the data centres were given the go-ahead?

Dr Bone Dodds: Heat networks, once built, are extremely flexible pieces of infrastructure. You can plug and play different heat sources. As you might be aware, the heat networks in Denmark started

off with a lot of coal and gas—heavy fossil fuels—but they are now on a rapid transition to other sources. Those networks are based on both a not-for-profit heat system and a competitive one, and the legislation says that they have to take the cheapest source of heat.

Waste heat tends to be the cheapest source of heat, which is why Lothian Heat exists. When we looked at the local heat and energy efficiency strategy—LHEES—for East Lothian, to start with, we saw that we had way more waste heat in the region than the residents of East Lothian could take. Edinburgh has the opposite problem: it has a huge density of homes and homes that are hard to heat, which need a lot of hot heat, not heat at lower temperatures. We thought that it would make sense to do as the Danes do and have big transmission pipes to take the heat from where it is to where it needs to be.

I do not have a strong position on what would be best for data centres. It would be extraordinarily difficult to use all of the heat from gigawatt-scale data centres. It might be a bit overwhelming to try to use all the heat from something so big. There would still be a huge amount of wasted heat.

I wonder whether Scotland and the UK deciding, on a speculative basis, when and where data centres get built by developers popping up left, right and centre like whack-a-mole is the most strategic approach. I suggested in my written submission that we should consider data centre proposals in alignment with regional energy strategic plans—RESPs—and that we should have proper heat and power plans, so that we look at a proposal in the round.

I have not examined whether there is any benefit to having smaller data centres as opposed to bigger ones, and I do not have any firm evidence on that right now. I can ask my colleagues in Lothian Heat whether there is any evidence. However, Octopus Energy has a partnership with a small data centre provider whereby the proposal is to build a small data centre, provide free heat for the village—I forget its name—and put in all the infrastructure for it. It will be on a much smaller scale than the gigawatt-scale data centres, but, apparently, the proposal passed through planning with almost no concerns from the residents, which is unheard of for a data centre proposal.

Decisions should be made strategically and in the widest interests of the economy, the environment, energy use and, ultimately, what the community needs and could truly benefit from. We should not build new sources of waste heat when people are going cold in their homes and we have a fuel poverty issue. Everybody of all ilks probably notices how difficult it is to pay their energy bills, regardless of whether they are in fuel poverty. It is a problem for all of us. Our big thing is that no heat should be wasted and no home should be cold.

Simon Kerr: If we had the heat network infrastructure that Denmark has, we could go large. We could put a data centre where we have access to oodles of electricity and it would not worsen constraint. Somewhere like Cockenzie would be perfect. In the early stages, as we grow our heat networks, smaller and more numerous data centres might be better. We could start to build the heat network on that basis, with a view to connecting to larger, heat transmission-scale networks.

We need to get to a heat transmission scale. I refer to running pipes from Dunbar into Edinburgh and decarbonising Edinburgh and everything in between by using all the heat in Dunbar and the heat and electricity that are available at Cockenzie.

We have known for some time that we need to get to that, but we are still talking about it. If we had that in place, it would be no problem to have a data centre as big as you like in Cockenzie and thermal storage as big as you like, because, at an industrial site, we would not have as much of a problem with nimbyism. There is another industrial site in Dunbar, and Cockenzie is an industrial site.

Those would be great locations for large-scale developments. Smaller-scale centres could work in the short term, as long as there was a vision, an understanding and a strategy based on their being connected to the larger, more strategic-scale heat networks that have massive seasonal thermal storage. Those networks can take and store heat all the way through the summer—they do not reject it into the already boiling atmosphere—and they then discharge it in the winter months. Not only will that help to balance the electrical grid, but there are the other benefits that come with scale.

We do need to get to scale. Smaller data centres might work in the short term, but, ideally, there will need to be multiple ones. They have the advantage of our not having to rely on any one heat source; in other words, if one were to disappear, it would not be the end of the world. That would not be the case if a large data centre were to disappear.

I asked our colleagues in Denmark about their experience with regard to the longevity of data centres, and they said that they last 20 to 25 years. They do not just stop—a lot of investment goes into them. Therefore, if you fear that they will just stop next year and move somewhere else, I have to tell you that that sort of thing has not been seen before. I am not saying that it cannot happen, though.

I would say that it is a case of either/or, but the fundamental thing is to get the pipes in the ground.

Line 1 of any planning application should highlight the heat reuse aspect; after all, it will solve the water issue, too. If you use heat network water to cool a data centre, there will be no evaporative cooling, so there will be no water wastage. In short, we should be thinking about the water issue, the just transition with the price of heat and the longevity and future proofing that a heat network can bring. If data centres can facilitate heat networks, I will be a big fan.

The Convener: You have touched on some of this already, but to what extent does the specific design of a data centre—that is, its source of electricity, its cooling system and so on—affect the generation of waste heat?

Simon Kerr: You are right—we have touched on this. A data centre must be designed for heat capture and reuse—that is the number 1 priority. If it is not, it should not be allowed anywhere near us.

We cannot allow heat to be dumped; we have to start valuing it. If it were food they were throwing away, there would be uproar about the people starving in our streets and so on. People are freezing, and we are allowing data centres to throw away heat—it is just not sensible. The just transition, net zero ambitions and the climate—everything points to the need to reuse heat. Data centres must not be allowed to dump heat in a modern society that is trying to do its best for people and the climate.

Moreover, this is an easy thing to do—it is not rocket science. We are talking about water going through a heat exchanger. The Romans did that to heat their baths. It is not technically difficult, so we should not let anybody get away with saying, “We can’t do it—it’s impossible.” It is pipes in the ground and a water-to-water heat exchanger to extract heat. It is dead easy—dead simple.

However, we do have to put the pipes in the ground. In Edinburgh, we are used to digging up our streets. I have lived here for 30 years and have lived through the trams, the gas pipe installations and all the rest of it. I know that there will be disruption, but we have to start doing it.

I cannot remember the rest of your question, but the infrastructure and the science behind heat reuse are dead easy and simple. It is old, not new, technology.

Dr Bone Dodds: In East Lothian, we carried out a feasibility study of what could happen to all the waste heat that exists—whether it could be used just in East Lothian and whether there would be any benefits in taking things to a bigger scale and bringing that heat into Edinburgh and beyond. Our plans were originally based on a seawater heat pump at the Cockenzie power station site, but, once the proposals for the data centre emerged, we were quite pragmatic and thought that, if we could capture that quite hot—perhaps 45ºC—heat, it could be a boost for the heat network. Our aim is to get sustainable heat at the cheapest price for people.

As for the design of the centres, it matters to us whether they will use fossil fuel energy and fossil fuel back-ups. If data centres do get permission and are allowed to use fossil fuels, it is worth noting that fossil fuel energy production itself gives off waste heat. It is imperative that we bear that in mind. It is not for us to say whether you should allow further use of fossil fuels for data centres, but we would think about whether the heat could come from green sources.

11:15

We have been trying to have conversations with data centre developers to find out what sort of systems they plan to use and whether they could interact with heat networks. Some conversations have seemed warm to begin with, but pinning them down to firm, up-front answers about what exactly they are planning to use, how they can interact with heat networks and how much they are willing to do to ensure that the heat is reused is not as easy as it should be when we think about the opportunity and about Simon Kerr’s comment that we should not be using a system if we cannot reuse the heat it produces.

We have not seen a strong approach so far, and the only way to make the approach stronger is to require it so that councils can ask for it in the heads of terms when they are talking to data centre developers. It has to be up front, as we cannot pretend that someone will care about it later if it is not there at the start. It is important to strengthen the approach.

There was a similar situation with Viridor Dunbar, which was a bit before my time. As part of the planning permission process, Viridor was asked to provide heat from its waste management plant, and the plant got planning permission because it said it would be able to provide that heat. Viridor might occasionally ask whether anyone wants that heat, but it is not likely that the plant will ever provide any heat. We cannot let the same thing happen with data centre developers. If it becomes material and part of what the Climate Action Committee or the Scottish Government decides makes a data centre worthwhile, there must be a mechanism in place to require that.

In my written submission, I offered some thoughts about how you might do that. You might ask a developer to contribute to a consortium of people who would figure out how to get some of those pipes in the ground. There is a lot of

momentum behind heat networks, but you need that momentum to build if things are actually going to become real. If a company that is building a large data centre says that it will help to put pipes in the ground, in order to get the heat from there to a major population centre, that will probably make heat network development happen more quickly. However, that will not happen on the basis of good wishes; it will happen only if we require it and create the infrastructure for it.

In Denmark, transmission pipes are owned by consortiums of municipalities that collaborate, with each municipality usually owning the distribution pipes that go from the main pipes into local places, but we do not have that system here. If we are really serious about having a transition to clean and just energy, and if it is to be one that incorporates what data centres can bring us and what heat networks can do to get us off the gas network, we must think about that strategically. We have to ask which sites can give us the best of all those worlds and then put effort into making a proper plan for how the heat will be used, so that it can be a catalyst for ramping up heat networks and so that there can be a realistic chance of developing them.

If data centres are going to be built, that is how it should be done, to ensure that they actually bring some benefits.

Donald MacKinnon: I will drill down into the idea of siting. Simon Kerr, you spoke about data centres as a gift from heaven for heat networks. How important is the siting and location of data centres in realising that? There is UK Government policy about zoning in North Lanarkshire, but siting is largely up to developers, who decide where they want to take projects forward.

Simon Kerr: It is very important. Proposals on heat network zoning are going through the UK Parliament, and there are plans for a heat in buildings bill in Scotland. We are crying out for that. As part of heat network zoning, you could have data centre zoning. To start with, you could look at where there is an abundance of electricity, and then at whether there is a heat network demand. Is it near a heat network zone? Is there electricity? If the answer to both questions is yes, that is a data centre zone.

We do not want to stick a data centre in the sticks or in the middle of nowhere, where it can do nothing with its heat and it will be allowed to dump it. There is heat network zoning, and we also have the regional energy strategic programme and the National Energy System Operator. Heat networks, hydrogen production and data centres are part of the regional energy strategic plan. That needs to co-ordinate with the heat network zoning programme, which needs to co-ordinate with the data centres. All of it needs to come together to identify strategic locations and strategic sites where green energy parks could be created. Those would be on sites where there is abundant energy available, whether or not that is electricity—there may still be a need for gas; who knows?

Denmark has plenty of examples of green industrial parks operating in symbiosis. There is a great example in which one company's waste heat is another's process heat, and so it goes on, until you are left with 80°C heat. As none of them wants that, they feed it into the local town, for its heating. There is a circular economy.

Data centres could tie into that, and that could tie in very nicely with the green industrial strategy, which we have talked about. You could say, "Okay—let's get the data centres into industrial parks," where no one will really mind too much if another big shed goes in that hums a little bit at night. It is in an industrial park; it is not in a neighbourhood where it will annoy people. That is part of the heat network zone—it is your energy centre and your large thermal storage. Your green hydrogen production could go in there, too, because around 30 per cent of it is high-temperature waste heat, which could also go into the heat network.

You have your green hydrogen, your data centre and your energy centre, and you can then bring in other industries that also produce waste heat. You can bring them all into that cluster and, all of a sudden, we have cracked it—we have more heat than we know what to do with. There are pipes in the ground and we can heat Glasgow and Edinburgh.

The location of the proposed data centres fits with the vision and strategy for where we need to get to by 2045. Interestingly, Denmark has just advanced its target for net zero, which was previously 2050. It has brought it forward by five years.

I do not know whether we are on track to meet our 2045 targets. I would say that we are not quite on track, but something like what I have described, with a vision and a strategy for helping out industry with waste heat, would be helpful. It would help industry on the environmental and social governance side, and it would help the investors, who would say, "Okay—you're a good guy; you're not dumping heat." All of a sudden, you are creating wealth and jobs, and you are cleaning the heat. It is, therefore, fundamentally important where the data centres, the heat energy centres and the hydrogen production go.

Donald MacKinnon: You touched on industrial uses for heat. In thinking about locations where home heating might be difficult to achieve, what

are the industrial uses for the heat from data centres? What are the big opportunities there?

Simon Kerr: It is too low for anything else. The old data centres, and some of the modern ones as well, are discharging heat at around 30°C—

Donald MacKinnon: So, it is not useful for any industrial processes.

Simon Kerr: No. For the heat network, we would increase the temperature. We would use the 30°C on a water-to-water heat pump, with water on one side and water on the heating side, both in a closed loop, so that no water was being discarded. We would use that 30°C as our source temperature and lift it to 70°C. The energy—the electrical input to a compressor—needed to lift the temperature by that amount would be a lot less than would be needed to lift the temperature from an air temperature of -5°C to 70°C, so you would be getting a lot more bang for your buck. Not only are the electrons—the energy—that you are putting into the data centre, through electricity, creating all the wonderful computer power that we all need; you are also now creating heat, so you are getting double usage out of those electrons.

The new data centres discharge heat at around 45°C, so they are getting hotter. Even 45°C heat is pretty useless for process heat in industry, but it is not bad for the heat network—again, going from 45°C to 70°C requires less electrical input, so you are reaching that temperature at an even cheaper price.

Yes, the location is very important, but I do not see an industrial use for the excess heat.

Donald MacKinnon: I put that question about the importance of the siting of data centres to Gemma Bone Dodds. It would also be helpful to understand more about the viable distances from the point of end use and the possibilities around that. At what point is it not really practical to transport that heat?

Dr Bone Dodds: The one thing that I would add to what Simon Kerr said is that, although I do not have any examples to hand, I have heard about locating lower-grade heat around smart agriculture—it might help with growing food, which will be a challenge for our climate going forward. Co-locating those could be something to pick up further.

The surprising thing in our feasibility study was that, although we think that it is a long way from Dunbar to Edinburgh, in Denmark it is completely normal to lay pipes of that size between places that are such distances apart. Larger transmission pipes are feasible; mostly, it depends on the land ownership and the usual questions about who has permission. If you are treated as a utility, that allows you to proceed without having to get all those permissions separately. So, the feasibility will depend on whether legislation makes heat networks a utility, as has been done for electricity and gas companies

Donald MacKinnon: I am sorry to interrupt, but is that distance the limit or the average? Where does it fit into what is possible?

Dr Bone Dodds: Again, it depends on the context. For the UK, that would be mind blowing, because we are so far behind. In other countries—

Donald MacKinnon: How about for Denmark?

Dr Bone Dodds: I would be happy to send the committee specific numbers about the longest pipes in Denmark.

Modern pipes are heavily insulated so that the heat loss between places is very small. The two issues are the cost of the pipes and their siting. If what you get for that initial infrastructure cost is a vast source of cheap or free heat, that would offset the cost of the pipe. It comes down to the business plans.

That is why we are saying that there needs to be a heat transmission plan. So far, our approach to heat networks has been to take on small ones. However, the area is not developing in that way. We might need to think about having more of a spine-and-spoke approach, which would fit with the idea that we should plan to put data centres where electricity is clean and freely available, and where it is feasible to send piping into places. EnergiRaven has done a lot of great work on the concept of heat highways.

We do that all the time with other industries, such as our water and our gas. We send gas, which is a combustible material, vast distances across the UK. This is about moving hot water around, so it is a lot less risky. One thing that got a lot of people excited when we did the East Lothian feasibility study was the plan for a 3m-wide, 30-mile long cycle route from Dunbar to Musselburgh. We could dig a bit deeper as the route was being built and stick the pipes underneath. Then, although there would not be much heat loss from the pipe, it might help with ice on the path in the winter.

We need to think smarter about how to do all that stuff, rather than having a whack-a-mole approach and thinking, “Wouldn’t it have been good if we’d designed it to be like that from the start?” That is what people want from us. Some of the pushback against net zero is from people seeing that things are not being co-ordinated enough and not getting the benefits up front. If we do this right, there is a real opportunity here.

Donald MacKinnon: Thank you. I have no further questions, convener.

David Barratt: I thank the witnesses for their input so far. Your enthusiasm is infectious—if we could bottle it and pour it into the Government, that would be fantastic.

I was struck by a few statements in Lothian Heat's written submission. It is clearly true that

"Heat that is available in principle is not the same as heat delivered".

I think that the convener covered the first bit of what I was going to ask. We have been talking a lot about the fact that the number of heat networks that are linked to data centres in the UK is very limited. A previous panel of witnesses told us about a 1MW data centre heating a small housing development or a pool—I cannot quite remember which it was. You have already touched on this, but is it practical for us to use all the waste heat from the data centres that are being proposed, given its magnitude, density and concentration? You talked about the gigawatt scale, but is it even practical to move and use all the heat from data centres operating at the 400MW, 500MW or 600MW scale?

11:30

Dr Bone Dodds: I will give you quite a rubbish answer again: it depends on whether we really put our backs behind heat networks, focus on getting them into place and take the opportunity, whether that comes from the waste heat of a data centre or from other sources—I refer to what Simon Kerr said about whether green hydrogen happens. It is about what catalyst makes that happen.

At the moment, we are expecting heat networks to create a whole new strategic infrastructure for the country, off their own back and largely driven by private industry. When that is done in a piecemeal way, we wonder why it is not happening faster. From Lothian Heat's perspective, I look at how we can do better with regard to fundamental human needs, such as the need to heat and power our homes. We have not valued heat in the infrastructure that we have put in in the past; we have discarded it and we continue to do so. Projects are taking place now that plan to throw it away, yet we stand around asking how we will solve fuel poverty. I get that it is not as simple as connecting those two things, because the plug and the connector are different, but we must find a way to work across those silos.

I have a bit of hope with the regional energy strategic plans—RESPs—which originally were only going to look at hydrogen and electricity, for some reason, not electricity and gas, which is our current energy supply as a whole. Those plans now include heat, and I imagine that data centres will also be on their radar.

There is a question as to how the planning process can take all of that into account. Getting a heat network up and running is not only about whether a data centre could be the catalyst for that but about whether the operator would be willing to put in some money to fund the pipes and create that catalyst for the quite complicated consortium of local authorities, private companies, community initiatives such as ours and contractors to come together around. With enough will, we can create wonderful things. However, the will to do that is perhaps spread out in too many places, so we need to concentrate it and try to prove that the model works. So, yes, I think that it is feasible.

David Barratt: I go back to the conversation that we had at the beginning. There is almost a catch-22 situation around whether there is scope to have a more dispersed, smaller-scale data centre network to input into heat networks when larger-scale data centres would be better able to afford the capital to build the pipes and the infrastructure that is needed to support such a big shift in the adoption of heat networks. Simon Kerr, what are your reflections on the first point of discussion and on the catch-22 situation around the choice between smaller-scale and hyperscale data centres?

Simon Kerr: To go back to your point, a 200MW data centre would generate a lot of heat, and the heat network market is maybe too immature for something of that size today. However, we should plan and look forward 10 years or so, which might be the length of time that it takes in the grid queue to get to that 200MW.

If we look at it from a heat point of view, rather than an electricity point of view, and go back to line 1 of the planning, we can consider how much heat we have and whether we have an off-taker for that heat. If the answer is that we do not yet have one but that we would like to provide that through heat networks, how do we get the data centres to contribute to that? How do we get the National Wealth Fund, the new energy wealth fund, regulatory asset base funding or whatever, in order to start strategically thinking about getting a suitably sized data centre for the scale of heat that will be required in the next few years?

That leads me to another thought. While the data centre is waiting five years in the queue, it will need back-up power, probably from air-source heat pumps installed at the start. During that time, we can get the heat network installed, use the air-source heat pumps, get everyone's gas boilers out and get the heat pipes in. In that way, when the data centre comes in, we will be ready—I could say "oven ready", but we should probably not use that phrase in this context. The heat off-take can then be taken from the data centre as soon as it is up and running.

The problem lies in what we have seen at Viridor and other energy-from-waste plants that have offered heat to whoever wanted it. North Lanarkshire Council was offered all the heat from DataVita's data centre but said, "No, thank you." East Lothian Council was offered the heat from Viridor's plant, which was great, and it had to be built with heat off-take, but the council did not want it.

We have to start saying yes to those offers, and the only way that we can do that is if we have large-scale heat networks. How will we get to that point? Everything comes back to whether we have the pipes in the ground. If we do, we can say, "Okay, we have capacity for a 20MW data centre in this area because that's the demand. We're going to grow, so, in 10 years' time, it'll be 50MW that we could put in, and, in another 10 years, it could be 100MW." It is about having joined-up, strategic, visionary thinking. One has a consequence for the other.

David Barratt: What requirements should be placed on the data centre developers in Scotland to ensure that they are ready to interact with heat networks? Helpfully, Lothian Heat's written submission has an eight-point plan that the Government could just pick up and run with. The Government is preparing guidance. Presumably, Gemma Bone Dodds, you think that all eight points should be adopted. Do you want to quickly touch on any of those requirements?

Dr Bone Dodds: Yes. All the points are about practical ways in which we can ensure that we do not end up with a Viridor-type situation. We are in a climate emergency; this country has been on fire and in drought this year and in previous years, and things will only get worse.

Our position is that we do not need data centres for heat—there are lots of other heat sources—but, if they get built, we want the heat to be used. If it is not used, we would not support that. That is why the definition of green data centres is important, as is requiring heat and power plans to be thought of up front. The feasibility and practicality of using the heat must be worked through actively by data centre developers instead of the heat being available and them smiling and doing nothing about it. We want to avoid that situation, with them getting permission and it being seen as a climate good but them never delivering on that good. We must manage that. The transmission plan, the strategic planning of energy, and consideration of data centres and heat in the round are the most important elements that sit behind it.

There is another thing that I want to say about heat networks. CopenHill is a really famous waste-to-energy plant that is just on the outskirts of Copenhagen.

David Barratt: That is the one with the ski slope.

Dr Bone Dodds: Yes. It is run by the municipality, which ran an international design competition to find a great look and use for it. On the hill is a ski slope, an 85m-climbing wall, a nature walk and a garden. It is also a big industrial use site.

Data centre developers are choosing to make these tin shacks, which people hate, so there is a wider planning point about who the developments are for and in whose interests they are being built. This might be a smaller issue, but data centre developers could learn from examples of best practice in relation to heat networks.

David Barratt: Simon Kerr, what requirements should we insist that developers meet?

Simon Kerr: Heat off-take must be considered. I do not know whether there is any mileage in providing data centres with some sort of financial encouragement. For example, if the heat is not being used, there could be a fine every year. A carrot-and-stick approach might be needed. That money could then be used to start putting pipes in the ground, and data centre developers would become active advocates of heat networks—they would demand heat networks and try to sell them to the local community.

I do not know whether you are familiar with the learning estate investment programme funding for schools in Scotland. That is brilliant funding. Schools are paid on the basis of their energy consumption. If they are below a certain benchmark, they get all their energy funded. The problem with Salix Finance funding and, previously, with private finance initiatives and public-private partnerships relating to schools was that people built whatever they could get away with. It did not matter whether what was built was inefficient. Those schools are now being handed back to councils, which cannot afford to run them. The new LEIP funding is brilliant, because it encourages developers to build good schools with heat efficiency in mind.

We should put in place planning policies with funding or penalties to encourage developers to build the best data centres with the best heat recovery and heat networks, and the amount of heat off-take should determine how little or how much they get fined. That could be a nice way to gently encourage those in the cash-rich data centre world. I do not believe that there is any poverty in data centre companies—they are cash rich and are making an absolute fortune—so let us use a sliver of that profit to tackle fuel poverty and decarbonise our heating systems.

I listened to the Economy, Tourism and Energy Committee's evidence session on 1 September

with Professor Dieter Helm and Dr Simon Gill, who said that fossil fuels still account for 77 per cent of Scotland's energy economy. We have done great work on electricity, but we have not touched heating and transport. It could be a blessing from heaven if we can use the data centre world to start chipping away at and making some inroads into the carbon percentage on the heating side. Otherwise, we will not achieve the 2045 target.

David Barratt: Providing a financial penalty for wasted heat is a fantastic idea, which I would love to build into our recommendations.

I have a final question. In the earlier evidence session, I mentioned that the BGS and the Mining Remediation Authority have recently undertaken a trial at the geoenergy observatory in Glasgow that involved injecting heated water into an abandoned mine. The initial test involved testing the thermal efficiency of putting water at 32 degrees into the mine and extracting it later. I understand that the next series of tests will use water at 40 and 45 degrees.

Is there scope for building that into our plans, given that the proposed data centres in the central belt would be built largely in areas that were previously mined, with mine works underneath that could be used? I do not know how the water could be injected into a mine in one area and extracted somewhere else or whether the mine works are, to some extent, the network that has already been built. I do not know what research has been done on that. Have you given much thought to extraction from mine works and to using them as a heat store?

Simon Kerr: That is an obvious option to explore. However, the problem with putting warm water into the ground is that it can start to dissolve minerals and rocks, so there is a limit to the temperature of the water that can be put into natural rock, depending on whether it is sandstone, a mine or whatever. There is also the issue of the contaminants that are already in the water, so a filtration process is needed. Therefore, things become quite complicated, but it is a really good idea.

Finland and Denmark are using underground thermal energy storage. The largest such site in the world, which has 1.1 million cubic metres of water, is in Finland. That will be able to heat a town with 30,000 dwellings on a heat network for a year without any additional heat source. That is the long-term potential of thermal storage. The water is not just in a mine; in effect, three large Thermos flasks have been put in to keep the water from being contaminated.

11:45

Using mine water is a great idea. However, thermal storage is quite cheap—that is not the problem. Geothermal heat is really interesting, particularly in Scotland and, apparently, in Northern Ireland as well, which are in the Goldilocks zone. The trouble is that geothermal storage needs a large heat demand, and it works better in some places than in others. Once the heat transmission system is in place, geothermal heat could be a fantastic additional heat source for the large heat networks. It is a chicken-and-egg problem; we need to get the pipes connected first, and then we can start.

Mine water provides a good theoretical study and it will be great if it works, but there are a lot of complications. Some of those complications may be unnecessary, because a thermal store is not exactly the expensive or difficult part of the whole equation. Do you remember the gas storage facility at Granton? We are talking about that sort of scale of thermal storage but for water, which can keep it warm all year round. We can put all our industrial heat into stores in the summer months and then feed it out in the winter months. That then comes back to the question of scale and where we locate data centres. Land is expensive in city centres, so the stores will need to be located on the outskirts of cities and in industrial parks.

We keep coming back to the same issue: once we have got the pipes in the ground and the right scale, other things will click into place, one of which could be using mine water.

David Barratt: Thank you. I do not know whether Gemma Bone Dodds wants to make any brief comments. I suspect that we are pushed for time, but I have no further questions.

Dr Bone Dodds: I will make a couple of points. I am a Geordie, and there is a lot more going on with mine working and mine heat, such as in Seaham garden village and Gateshead Council, where I used to work. Gateshead Council put in heat networks before they were cool, because it made sense to decarbonise all the council buildings and high-rise flats that are in the town centre. It did that by using an innovative mechanism with a gas combined heat and power plant that played the flexibility market, which enabled the council to provide cheap heat. That has now been shifted to a mine water heat system. Mine water heating is really interesting.

On the point about thermal storage, a big opportunity for Lothian Heat in East Lothian is the Tarmac cement plant. That could, in itself, be a massive source of heat—way better than a data centre—if the heat could be captured. The plant also has a massive pit in the ground. When I had a conversation with the company, I was told that,

when it comes to decommissioning, if we would prefer the pit not to just be filled back in with soil and instead be tanked out to become a massive interseasonal thermal store, somebody just needs to tell the company that in advance of the final investment decision and tell it how that will be paid for, what the mechanism will be and who will own the pit going forward. That feels like the kind of thing that should be simple, because it clearly should work.

Having a big thermal store can stop a lot of the wind curtailment that is happening, because that is where a lot of the offshore wind in Scotland comes on board. Instead of paying for wind farms to turn off, the hot water can be heated up, get tanked and stay there. It would also help with curtailment more than a data centre, because wind blows the most in the winter, which is when we need the most heat.

We have been trying to make arguments about having a mechanism for flexibility. We have been trying to work with SP Energy Networks to do a study on how much a transmission-scale heat network and a large thermal store would save in curtailment payments, and how much would be saved by reducing the amount of grid upgrades that are necessary. That will be a huge cause of rising energy bills in the future, and it would be good if a heat network could prevent that.

Perhaps I can highlight another thing that a heat network does. If you have a lot of air-source heat pumps, they will probably turn on at the same time, which will increase the peak use of the grid. As a result, what we are building out with these heat pumps is not big baseload, but something to meet those big, sharp peaks. A heat network will flatten such a peak, so you do not need as many grid upgrades. As a result, your electricity and heating bills will be cheaper. If we are talking about the interaction between the sort of thermal storage that I am interested in and data centres, we could put a big data centre on the Torness site when it goes, because there would be a big thermal store next door. I am not saying that that should happen, and, of course, we will not know until we take a strategic look with that overarching plan.

David Barratt: Thank you.

Sanne Dijkstra-Downie: You have talked a little bit about sites, particularly in relation to distance from the heat source. Is there anything else that we should be thinking about in terms of suitable sites for these data centres, particularly with regard to the possible use of waste heat?

Dr Bone Dodds: It is not just about the heat—I have been reading a few articles lately about how the security risks posed by data centres need to be considered, too. With the military becoming more AI-enabled, centres have become strategic targets in the Iran war and in the Russia-Ukraine war, simply because so much activity gets run through them. It might pose a material risk for heat networks if military operations were to be run through data centres, but that is probably the only other thing that I can think of with regard to sites. It goes back to the need for a strategic plan for these places, so that we can think about security as well as other issues, but I have probably said all that I want to say about that.

Simon Kerr: The siting of data centres or any industry producing waste heat becomes an awful lot easier when you have heat networks of the scale that allows you to have multiple locations. If you were to superimpose the heat network in Copenhagen, for example, on the central belt, it would connect Glasgow, Edinburgh and Stirling. That network is now being expanded to Kalundborg, where they make the slimming drug. There are other industrial waste heat sources in the city, and those who are putting the data centre into Kalundborg are part-funding the 30-mile pipe to the nearest connection to Copenhagen's heat network. Data centres are desperate for locations that they can go to freely, even if it means having to pay for a bit of pipework. That is what is happening in Denmark, so we should not let them off the hook in this country.

The beauty of that 30-mile stretch is that it takes in three gas towns. Denmark still has gas: 70 per cent of the country is on a heat network, with maybe 10 per cent on gas and the other 20 per cent on heat pumps, and there are three towns between Kalundborg and Copenhagen that are still on gas. All of the waste heat going from Kalundborg to Copenhagen will allow the heating in those towns to be decarbonised. That stretch would be very similar to the stretch between Dunbar and Edinburgh, with Haddington, Dalkeith, Macmerry, Cockenzie, Prestonpans and Musselburgh on the way in.

Location does matter, but we need the pipes to go in.

Sanne Dijkstra-Downie: I was just going to ask you whether you think that data centre developers should be asked to contribute to the cost of the pipes, but I think that you have answered that question. Gemma, do you agree?

Dr Bone Dodds: Yes. I think that, if data centres get approved alongside some vague and unspecified community benefit scheme, there is a risk that the money will have no real impact. We need to think about requirements and how we get over the humps.

I know that heat networks feel really difficult now, but that is because we do not have them. Moreover, I think that we have lost something when it comes to doing big infrastructure projects

in the UK. To use the example of a slightly different mechanism in Denmark, By & Havn—which means "ports and harbours"—is a publicly owned development company for the city of Copenhagen that takes sites and develops them strategically. It sets out exactly what it wants to see on a site: the economic use, the energy use and the trees, biodiversity and nature. It does all the planning and creates a master plan, and then it sells off the site to the people who want to provide that. In doing so, it pays for the infrastructure that goes in.

Last year, I visited a site where the developers have been reclaiming more of the land from the sea. Basically, they had a problem in that, when they were making new things, they had a lot of waste soil, so they are putting it in the sea to create new land, and they are developing that land. The city council then has the money to pay for the metro to move on to the next place and serve it.

There is something in that story about the state accepting its powers of planning—what it can achieve and the markets that it can create through planning. Rather than thinking simply, "Oh, we're just trying to stop them doing a bad thing," the state is actively making a good thing happen. We are quite far away from that model in the UK, but that does not mean that it is not something that we should aim for. At the very least, if these companies are getting money—because they are there to make money—and we are not asking them specifically to ring fence some of that money for some of the priorities that we know the Scottish Government cannot pay for entirely by itself, that seems like it would be a logical way for us to consider trying to create revenue for the capital expenditure that is needed.

Sanne Dijkstra-Downie: Thank you. I will move on to a completely different question, on a tricky area that we have discussed with previous panels, which is retrofitting and the feasibility of having heat networks installed in areas where there are existing properties and buildings. As you said, Gemma, retrofitting is always harder to do, but we have heat network zones, as determined in LHEES, which include a lot of existing properties. My worry is that if we think about those as opportunities that work only for buildings that are still to be built, we lose opportunities there. Can you tell us a bit about how you see the feasibility of retrofitting heat networks in existing areas? Is it possible but just difficult and expensive? Can it be done, or is it really not something that is going to work?

Dr Bone Dodds: Lothian Heat exists, in part, because, while there are some great heat network developers—and I do not necessarily want to knock any of them in particular—they have mostly chosen to focus on new-build properties and connecting those to heat networks, not on retrofitting houses.

I want to clarify something. When I said earlier that retrofitting is difficult, I meant with regard to retrofitting the heat capture to existing data centres, rather than houses. Sorry—I could have made that clearer. In fact, heat networks—again, I think that air-source heat pumps are fabulous; this is not a competition, and we need all these technologies—are really simple to put in. I highly recommend that any of you go to the Energy Training Academy in Dalkeith, where you can see the heat interface units. I am sure that Simon Kerr could organise a visit.

Basically, it is a four-hour job. Once the pipes are in the road outside, you take the pipe from the road down to wherever your boiler currently is. You take the boiler and the gas pipes out, and you put the two pipes in—one heat pipe going in with the hot water and one heat pipe coming out with the cool water, once it has been round your house. That goes into the heat interface unit, which is about the same size as a boiler and is put in where your boiler would go. That is it. You might need to change some of the thermostatic radiator valves on your radiators to a different type, but, in terms of the impact on the home owner, it is probably one of the simplest technologies to install in switching over to clean heating. The other benefit of heat networks is that heat interface units are much simpler to use. Because they are not individual little gas power plants, they are safer and they last a lot longer than a gas boiler would, so the replacement cost is quite small. Lothian Heat's model, on which the feasibility study is built, is trying to get over the retrofit problem where the private sector has said that retrofitting is too much hassle.

I went to Aberdeen to visit Aberdeen Heat and Power, which is a fantastic publicly owned company that runs what is probably the largest heat network in Scotland. It was connecting the council housing stock to an expansion of the heat network, but it was bypassing the houses of people who had exercised the right to buy because it did not have a model for how to connect those people, if you see what I mean. We have not yet cracked the issue.

12:00

This year, Lothian Heat supported three applications to the heat network support fund. One involved connecting Musselburgh to the existing heat network at Millerhill. The fact that there is already a heat network there lowers the risk. We wanted to know whether we could connect Queen Margaret University to that network. It wants to decarbonise its heating, because it has some big gas boilers, and we wanted to connect about 2,000

homes in Musselburgh specifically to test the idea of retrofitting. Unfortunately, we were knocked back for feasibility funding. We were asked, "What proof do you have that people want this?" That was the point of the feasibility funding. When it comes to what we are trying to do on the ground, the bar is being set too high. We have the people on the ground who are ready to do that work.

The other application was for Porty Community Energy, which, as you will know, is a fantastic community organisation in Edinburgh. It has received 10 years of funding from the Regenerative Futures Fund to develop heat networks for the city, but its application was rejected for similar reasons. We have been making the case that we need to figure out what offer regular people will jump at when a heat network comes to their area. That is not necessarily the way that private companies look at the issue, but, because we are a community initiative, we want to figure out what people need, what they would accept and what good would look like to them.

We want to copy the Danish model as much as we can and bring that here. In essence, that model involves the provision of heat as a service. People pay for their heat, but they do not pay 15 grand for the installation and they do not pay when their boiler breaks. They ring up the heat network, which will come along to fix it. We are trying to take heat off people's list of things to be stressed about. Not-for-profit heat networks make it possible to do that kind of work, because, at the end of the day, they simply use the money to do the right thing for the customer.

Sanne Dijkstra-Downie: Would you say that the not-for-profit model could work in the case of data centres?

Dr Bone Dodds: For data centres?

Sanne Dijkstra-Downie: I am thinking about whether the waste heat from data centres could be used as part of the not-for-profit model.

Dr Bone Dodds: We would need to look at that. In Scotland, we have learned a lot about the Danish model, but we have been a bit reticent about the Danes' economic and business model, which is not for profit but also competitive. The Netherlands did what Scotland has been trying to do, which was to have quite a high level of private ownership of heat networks. Many of those networks are being brought back into public ownership because, in many ways, that approach has failed. We need to think seriously about expanding the range of ownership options in the long term, because we are advocating for community wealth building rather than building the same broken system of the past.

Sanne Dijkstra-Downie: Simon, do you have anything to add on retrofitting or business models that you think would work well here?

The Convener: We will move on after that question has been answered.

Sanne Dijkstra-Downie: Yes, of course—that was my last question.

Simon Kerr: We have close ties with Denmark, where retrofitting is being done. We work with Danish companies, so we see how they do it in Denmark.

Currently, 70 per cent of Danish dwellings are connected to heat networks. By 2030—in four or, rather, three and a bit years—80 per cent of them will be connected to heat networks. That is the scale of what is happening in Denmark, which is why the Danes have brought forward their net zero target from 2050 to 2045. Here, about 1.5 to 2 per cent of dwellings are connected to heat networks, so we have a mammoth task in front of us. Waste energy from data centres could be the catalyst. There is waste heat out there that is so cheap that it is being thrown away. We must use that to expand our heat networks.

A lot of our Danish colleagues have little summer houses on their islands. They have a lovely life in Denmark. They have heat pumps in their summer houses and they live on a heat network in their city houses. If you ask them what they would rather live with, they will all say that it is easier to live on a heat network than with an individual heat pump, because it is up to you to fix that individual heat pump should something go wrong. With the heat network, it is just heat off the street; it is a utility, like gas, water or electricity. It is easy to live with. That is the practical route that we need to take when it comes to the just transition.

The Scottish Parliament building cannot have heat pumps. This building's only route to decarbonising is to connect to a heat network. If that does not accelerate the push for heat networks, I do not know what will. Once you have those big anchor loads, you can start going into Dumbiedykes and all the rest of the fuel-poor areas. That is at a not-for-profit level, because no private company will spend hundreds or thousands of millions of pounds putting in tiny connections to flats that hardly use any energy and where the bills are very small. That whole funding model—a private company seeking to get dividends to make the return on its investment—does not exist in this context. That is why the Lothian Heat model of being not for profit and ploughing everything back in is the only way to tackle fuel poverty and fulfil the just transition part of it.

Ariane Burgess: We have been talking very excitedly about waste heat and the importance of capturing it. It has been a very interesting session, and it will help us with other aspects of our work.

I want to bring us back to hyperscale data centres and go into a space that you might not be across but that you might have come across, which is the heat island effect. I will ask the question but will pull back if you have not come across it or are not working with it. To what extent would hyperscale data centres result in a heat island effect if we did not do all the amazing stuff that you have talked about in terms of waste heat being captured?

Dr Bone Dodds: We have not looked at that yet, but I listened to a podcast this morning, on Radio 4, that looked at Slough and the heat island effect that is being seen because of the many data centres that are located there. I do not have information on the issue, but that might be an avenue that you could find something from.

Simon Kerr: Current and old data centres are throwing away 30°C heat, which is an awful lot of heat going into the atmosphere. If a data centre is hyperscale, it could produce 100MW at 30°C that goes into the local area during a hot summer.

As we talked about earlier, the problem is that data centres' heat rejection levels are getting higher—it is higher-grade heat. The AI ones are 45°C, which means that we are now dumping 100MW of 45°C heat into a local area in the middle of a heatwave. Regardless of whether that is the right or just thing to do environmentally—it clearly is not—we cannot allow heat to be continuously dumped on people from a health and quality-of-life perspective.

Ariane Burgess: How well developed is the research on the potential heat island effect?

Simon Kerr: I do not know, but I imagine that there are a lot of studies on it.

Ariane Burgess: Earlier, you talked about not understanding the value of all the different sources of heat waste. In considering whether AI data centres should be built in certain places, what other sources of heat island effects should we be aware of?

Simon Kerr: Grangemouth is a good example. The amount of waste heat that Grangemouth throws away every year is enough to heat Edinburgh. Nobody talks about it. Nobody says anything. You then have to take into account flaring, which is environmentally friendly. We flare methane to create CO_2, because methane is worse than CO_2, but, if we are going to burn methane, let us turn it into electricity and use it in a heat network.

We have to wake up and stop being silly about all the energy that we are throwing away. That includes curtailed wind, for which the bill is getting bigger: it is £2 billion this year, and it will get bigger still. The energy price cap has cost the UK Government £41 billion since 2021, and we are in the middle of that—it is only going to get worse. We need to wake up and start having some commonsense conversations about what we do with the heat, the energy and the wind that we have available, and use it a bit more sensibly.

The bit that ties it all together is heat networks—simple pipes in the ground that carry water.

Ariane Burgess: In the interests of time, I will make just one comment myself: I believe that the Parliament is part of a forthcoming heat network project, so I look forward to this building being an anchor point.

I put this question to the previous panel; it is specifically on heat networks connected to hyperscale data centres. From the work that you have been doing, are you aware of other countries bringing forward policy on hyperscale data centres that you think we should have a look at?

Dr Bone Dodds: I asked our colleagues in Denmark about that and whether there was more that they could share with us, but that information is not publicly available at the moment. Even though Denmark has the heat network infrastructures ready, there is still a very live and hot political debate, as there is in the UK, with regard to the acceptability of data centres and the rationale for why we would want to build them. Denmark is facing exactly the same challenges as we are. The heat question has been answered to a greater extent there, but the fact that Denmark is still challenged by the question of the acceptability of data centres is interesting.

Ariane Burgess: So, Denmark would be a country worth checking in with.

Dr Bone Dodds: Yes, absolutely.

Ariane Burgess: Great. Thank you.

Simon, do you have anything to add on that?

Simon Kerr: I would highlight the Scandinavian countries in general, including Sweden and Finland, and Germany, which has a good progressive approach. Germany was hit quite badly in respect of the Russian gas supply and stuff like that, so it has had to transition. The Netherlands faces a similar situation to us in that it has a large gas network that it is decarbonising and trying to get off. It has adopted the three Danish heat network laws and is pushing forward quite fast on its heat network development, so we could look to the Netherlands and see what its

plans are for data centres. I do not know whether it has any, but I would be surprised if it did not.

We can look at what Germany, the Netherlands and the Scandinavians are doing. Northern Europe is where data centres want to be, and that is where heat is required, so, again, that could be a very nice match. I would stick to the north—I would not go down to Arizona and places like that, because they would just dump the heat because they have no need for it, and they would use evaporative cooling and a lot of water. That would just be silly for countries that need heat.

Dr Bone Dodds: The example from Viegand Maagøe is in our written submission, at paragraph 129. It is also worth highlighting that Germany's Energy Efficiency Act requires new data centres to achieve an energy reuse factor of at least 10 per cent, rising to 15 per cent and then 20 per cent. Germany is saying, "If you're using 100MW, you must reuse that"—presumably in the form of heat, or they might find other ways to do it—and it is ramping up the percentage over time. David Barratt asked about other mechanisms, and that is an example of a carrot-and-stick reuse policy.

Martyn Day: What can be done to mitigate or lessen the impact of heat escaping from data centres? Simon, perhaps you would like to start.

Simon Kerr: I would guess that, if you build a leaky data centre, you might not need to cool it as much. However, data centres are very secure places; they are very protective and there is barbed wire everywhere—you cannot get anywhere near them. I am guessing that they are going to be built, or they should be built, to a high standard.

The heat escaping from a data centre is not a problem that I have heard of as being a concern. I am not an expert on data centres by a long way, but, once you have the heat exchanger in place and the water cooling, that is it. It may be different with air cooling—if you have air cooling on the racks, blowing air across, I am sure that that air will find a way out. If you had a wet system—a closed-loop system with water in it—it would be the same as the air-conditioning system in a hospital or an office block, and we do not get leaks there.

It depends on how the data centre is cooled. If it is cooled with water, I do not think that it is an issue, whereas with air it may be.

12:15

Martyn Day: Following on from that, what should our planning policy set as an exceptional level for heat emissions?

Simon Kerr: Zero into the air—no, that may be a bit extreme, but there really is no reason why the vast majority of the heat that is taken in order to cool the racks could not be transferred and used. Heat exchangers have about 95, 98 or 99 per cent efficiency, so they can transfer that heat very easily into a heat pump and then lift it up. The level should be minimal; there really should be very little heat coming out of data centres.

Dr Bone Dodds: You refer to heat escaping, but that is not the case. A data centre will either be pushing the heat out on purpose or it will have a closed-loop system and the company will be dumping that heat into some water—probably the sea or something like that—or maybe into the air. That happens by design, so the issue is about requiring the right kind of design up front.

From what I have read, it seems that data centres are moving more towards closed-loop technology because of the pressure around water and the efficiency of the system. Essentially, however, unless companies are made to care about what happens to that energy once they are done with it, we should not expect them to care innately. They want the data to run the processors, and, unless they are made to care, that is probably less of an interest for them. Just requiring the right kind of closed-loop systems, and connecting with heat networks and reusing that heat, would be the way to stop heat escaping.

Martyn Day: Thank you. That is helpful.

The Convener: I thank our witnesses. There has certainly been a great deal of information presented to the committee, and I thank you for that.

That concludes the committee's business in public, and, as agreed earlier, we will move into private session to consider the remaining items on our agenda.

12:17

Meeting continued in private until 12:55.

This is a draft *Official Report* and is subject to correction between publication and archiving, which will take place no later than 35 working days after the date of the meeting. The most up-to-date version is available here: https://www.parliament.scot/chamber-and-committees/official-report

Members and other meeting participants who wish to suggest corrections to their contributions should contact the Official Report.

Official Report
Room T2.20
Scottish Parliament
Edinburgh
EH99 1SP

Email: official.report@parliament.scot
Telephone: 0131 348 5447

The deadline for corrections to this edition is 20 working days after the date of publication.

Published in Edinburgh by the Scottish Parliamentary Corporate Body, the Scottish Parliament, Edinburgh, EH99 1SP

All documents are available on
the Scottish Parliament website at:

www.parliament.scot

Information on non-endorsed print suppliers
is available here:

www.parliament.scot/documents

For information on the Scottish Parliament contact
Public Information on:

Telephone: 0131 348 5000
Textphone: 0800 092 7100
Email: sp.info@parliament.scot



The Scottish Parliament
Pàrlamaid na h-Alba